



Publication Year	2021
Acceptance in OA	2025-03-08T18:03:17Z
Title	SOFIA 2 - an automated, parallel H I source finding pipeline for the WALLABY survey
Authors	Westmeier, T., Kitaeff, S., Pallot, D., SERRA, Paolo, van der Hulst, J. M., Jurek, R. J., Elagali, A., For, B. -Q., KLEINER, DANE, Koribalski, B. S., Lee-Waddell, K., Mould, J. R., Reynolds, T. N., Rhee, J., Staveley-Smith, L.
Publisher's version (DOI)	10.1093/mnras/stab1881
Handle	http://hdl.handle.net/20.500.12386/36552
Journal	MONTHLY NOTICES OF THE ROYAL ASTRONOMICAL SOCIETY
Volume	506

SOFIA 2 – an automated, parallel H I source finding pipeline for the WALLABY survey

T. Westmeier^{1,2,★}, S. Kitaeff^{1,3}, D. Pallot^{1,3}, P. Serra⁴, J. M. van der Hulst⁵, R. J. Jurek⁶, A. Elagali^{7,2},
B.-Q. For^{1,2}, D. Kleiner^{1,4}, B. S. Koribalski^{1,7,8}, K. Lee-Waddell^{1,7}, J. R. Mould⁹, T. N. Reynolds^{1,2},
J. Rhee^{1,2} and L. Staveley-Smith^{1,2}

¹International Centre for Radio Astronomy Research (ICRAR), M468, The University of Western Australia, 35 Stirling Highway, Crawley WA 6009, Australia

²ARC Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D), Australia

³Australian SKA Regional Centre (AusSRC), Australia

⁴INAF – Osservatorio Astronomico di Cagliari, Via della Scienza 5, I-09047 Selargius, CA, Italy

⁵Kapteyn Astronomical Institute, University of Groningen, Postbus 800, NL-9700AV Groningen, the Netherlands

⁶Department of Physics, School of Mathematics and Physics, The University of Queensland, Brisbane QLD 4072, Australia

⁷CSIRO Astronomy and Space Science, Australia Telescope National Facility, PO Box 76, Epping NSW 1710, Australia

⁸Western Sydney University, Locked Bag 1797, Penrith NSW 2751, Australia

⁹CAS, Swinburne University, PO Box 218, Hawthorn VIC 3122, Australia

Accepted 2021 June 29. Received 2021 June 28; in original form 2020 September 29

ABSTRACT

We present SOFIA 2, the fully automated 3D source finding pipeline for the WALLABY extragalactic H I survey with the Australian SKA Pathfinder (ASKAP). SOFIA 2 is a reimplement of parts of the original SOFIA pipeline in the C programming language and makes use of OpenMP for multithreading of the most time-critical algorithms. In addition, we have developed a parallel framework called SOFIA-X that allows the processing of large data cubes to be split across multiple computing nodes. As a result of these efforts, SOFIA 2 is substantially faster and comes with a much reduced memory footprint compared to its predecessor, thus allowing the large WALLABY data volumes of hundreds of gigabytes of imaging data per epoch to be processed in real time. The source code has been made publicly available to the entire community under an open-source licence. Performance tests using mock galaxies injected into genuine ASKAP data suggest that in the absence of significant imaging artefacts SOFIA 2 is capable of achieving near-100 per cent completeness and reliability above an integrated signal-to-noise ratio (SNR) of about 5–6. We also demonstrate that SOFIA 2 generally recovers the location, integrated flux, and w_{20} line width of galaxies with high accuracy. Other parameters, including the peak flux density and w_{50} line width, are more strongly biased due to the influence of the noise on the measurement. In addition, very faint galaxies below an integrated SNR of about 10 may get broken up into multiple components, thus requiring a strategy to identify fragmented sources and ensure that they do not affect the integrity of any scientific analysis based on the SOFIA 2 output.

Key words: methods: data analysis – software: data analysis.

1 INTRODUCTION

Several precursors and pathfinders to the Square Kilometre Array (SKA; Dewdney et al. 2009) have recently begun taking early data, including the Australian SKA Pathfinder (ASKAP; Hotan et al. 2021), MeerKAT (Jonas 2009; Camilo 2018), and the APERTure Tile In Focus (APERTIF; Verheijen et al. 2008). One of the major surveys to be carried out with ASKAP is the Widefield ASKAP L-band Legacy All-sky Blind Survey (WALLABY; Koribalski et al. 2020) that is expected to image the entire sky south of a declination of $+30^\circ$ in the 21-cm line emission of neutral hydrogen (H I). WALLABY is expected to produce about 1 petabyte of H I imaging data and detect the H I emission from approximately half a million galaxies out to a redshift of $z \approx 0.26$.

Given the unprecedented amount of imaging data anticipated from extragalactic H I surveys such as WALLABY, the detection and characterization of galaxies will need to occur in a fully automated fashion with minimal manual intervention using dedicated spectral-line source finding software such as DUCHAMP (Whiting 2012), SÉLAVY (Whiting & Humphreys 2012), or the Source Finding Application (SOFIA; Serra et al. 2015). The performance of several source finding packages and algorithms, and their suitability for H I source finding, was tested by Popping et al. (2012).

SOFIA has been specifically developed for the purpose of detecting galaxies in extragalactic H I surveys like WALLABY. It features several powerful algorithms that have greatly improved the quality and accuracy of H I source finding (e.g. de Blok et al. 2018; For et al. 2019; Blue Bird et al. 2020), including the *smooth and clip* (S + C) algorithm and a new method for automatically identifying and removing unreliable detections (Serra, Jurek & Flöer 2012a).

* E-mail: tobias.westmeier@uwa.edu.au

While SOFIA’s algorithms are well-suited to detecting galaxies in WALLABY data, there are several shortcomings to the current implementation of SOFIA that make its application to large data volumes challenging, most notably its large memory footprint and comparatively low speed. In order to address these and other issues, we have reimplemented the most powerful algorithms of SOFIA in the C programming language in order to create a much faster and more memory-efficient pipeline called SOFIA 2. In addition to improvements in the implementation, we employ techniques such as multithreading and parallelization to significantly speed up the processing of large data volumes, allowing SOFIA 2 to process a single 800 GB WALLABY data cube in a matter of minutes rather than hours on a modest number of computing nodes.

Here, we introduce SOFIA 2 and present its basic performance measures based on tests with both mock data and real WALLABY H I data. A brief overview of the implementation of SOFIA 2 and the main improvements compared to the original SOFIA (hereafter referred to as SOFIA 1 for clarity) is presented in Section 2, followed in Section 3 by a description of the parallel framework named SOFIA-X. Section 4 outlines some of the core algorithms employed by SOFIA 2. In Section 5, we present tests of the speed and memory usage of the pipeline, while the performance of SOFIA 2 with respect to completeness, reliability, and parametrization accuracy is presented in Section 6 based on tests with mock galaxies injected into genuine ASKAP H I data. Initial results from the processing of WALLABY pilot survey data with SOFIA 2 are presented and discussed in Section 7 followed by a summary and more general discussion in Section 8.

2 IMPLEMENTATION

SOFIA 2 is a reimplement of the core algorithms of the original SOFIA 1 pipeline, most notably automatic flagging, noise normalization, S + C source finder, and reliability estimation (Serra et al. 2012a). Some of these algorithms have also been simplified and improved compared to the original SOFIA implementation, while other, less widely used algorithms and features have not yet been ported, including the 2D–1D wavelet filter (Flöer & Winkel 2012), the CNHI source finder (Jurek 2012), position–velocity diagrams, automatic scaling of the reliability measurement kernel, and more sophisticated options for controlling the shape and size of the smoothing kernels in the S + C finder.

Unlike the original SOFIA 1 pipeline, which was written in a combination of PYTHON, CYTHON, and C++, SOFIA 2 is written entirely in the C programming language, more specifically using the C99 standard (formally known as ISO/IEC 9899:1999). This has resulted in several substantial improvements:

- (i) The conversion to C alone has resulted in a modest improvement in speed and a substantial improvement in memory usage as compared to SOFIA 1, even without multithreading or parallelization (see Section 5 for details).
- (ii) As time-critical algorithms can be coded directly in C, the number of external library dependences has been reduced to just one (WCSLIB; Calabretta 2011), making SOFIA 2 much easier to install and maintain.
- (iii) Thanks to the native support of OpenMP by commonly used C compilers, most time-critical algorithms in SOFIA 2 have been fully multithreaded, resulting in additional significant gains in speed on multicore architectures.

The choice of OpenMP (Dagum & Menon 1998) for multithreading was made to keep the code as simple and easy to use as possible. OpenMP is natively supported by many compilers and therefore does

not impose any additional dependences. In addition, it can be adopted with just minimal modifications to the source code and can easily be disabled if necessary.

Like its predecessor, the source code of SOFIA 2 has been publicly released on GitHub¹ under an open-source licence to make it freely available to the entire astronomy community. While the code has been optimized for compilation with the GNU Compiler Collection’s GCC compiler,² it should in principle compile with any C compiler that supports the C99 standard and is therefore expected to be able to run on a wide range of POSIX-compliant systems.

As the main purpose of SOFIA 2 is the automated detection and parametrization of galaxies in the large volumes of data produced by H I surveys such as WALLABY, we do not provide a graphical user interface as with SOFIA 1. However, like its predecessor, SOFIA 2 can be easily controlled through basic parameter files and comes with a robust set of default parameter settings that can serve as a starting point for establishing optimal parameter settings.

3 PARALLEL FRAMEWORK

A single WALLABY spectral image cube is expected to be up to 800 GB in size, which immediately raises two issues. First, it would take a significant amount of time to process such a large data volume on a single computing node even with multithreading enabled. Secondly, up to 2 TB of memory would be required to load and process such a cube all at once (see Section 5). The only viable solution to these two problems would be to split up the data cube into multiple subregions each of which must be small enough to fit into the amount of memory available on a single node. The individual subregions can then be processed in parallel on separate computing nodes and the resulting catalogues be merged at the end of this process into a single source catalogue of the full region.

SOFIA 2 already has the capability to read arbitrary sections of a data cube. The merging of multiple outputs, however, would still need to be carried out manually, including the difficult task of resolving cases of duplicate detections in overlapping regions. In order to automate this entire process, including the merging of output from multiple, partially overlapping regions of a data cube, or from multiple runs of the source finder on the same region, we have developed the SOFIA-X framework.

SOFIA-X is a wrapper around SOFIA 2 that can spawn multiple instances of the pipeline across different nodes of a computing cluster using the Slurm Workload Manager (Yoo, Jette & Grondona 2003) and collate the output from all instances into a single, coherent source catalogue. The choice of Slurm was again driven by the idea of simplicity and ease of use, as Slurm is open-source and readily available on most supercomputers. The instances of SOFIA 2 processing the regions require no interprocess communications and therefore can be executed as an array job in Slurm. The scripts can be easily changed if a different batch job scheduler is desired or available.

Each instance of the source finding pipeline is provided with a configuration file that specifies a region of the data cube to be read and produces a catalogue of detections for that region in VOTable³ format. There are no limitations on how these regions are defined, and it is entirely up to the user to choose suitably sized spatial and/or spectral regions with adequate overlap by using the `input.region` keyword in the SOFIA 2 parameter file.

¹<https://github.com/SoFiA-Admin/SoFiA-2/>

²<https://gcc.gnu.org/>

³<http://www.ivoa.net/documents/VOTable/>




Figure 1. SOFIA-X data base schema showing the field names and data types in each of the four tables used. Each *run* of SOFIA-X spawns multiple parallel *instances* of SOFIA 2 that will typically operate on different regions of the data cube. The source parameters for each *detection* will be written to a common table, with another table storing the image *products* associated with each detection.

Once the VOTable is created, it is then passed to a PYTHON script that imports the catalogues and image products for each detected source into a data base on a server. The location of the data base server can be the same for all runs of SOFIA 2 on a cluster, or the scripts can be pointed to the different locations of multiple data bases. Each run and each instance of the programme are uniquely identified in the data base. The SOFIA 2 parameter settings of the runs are also recorded in the data base as metadata to enable provenance of the data and retrospective reviews. The data base schema is presented in Fig. 1.

During the import of each catalogue, a PYTHON script attempts to automatically identify and remove duplicates in overlapping regions by checking if the positions and a few other observational parameters agree within certain thresholds specified by the user. If the heuristics of the automated decision making does not allow unambiguous resolution of duplicates, then the detections will be flagged for manual resolution in the data base. For example, this could be the case if the positions of two sources agreed within the specified thresholds but their fluxes did not.

SOFIA-X provides the user with a convenient web portal for inspecting the merged catalogue and manually resolving any flagged detections. The portal has a *Table Access Protocol* (TAP)⁴ interface that allows users to directly connect through VO-compliant software,

such as TOPCAT (Taylor 2005), or run ADQL⁵ queries on the source data base. The image products of individual sources are accessible through VO-compliant data links and can be directly visualized with software such as ALADIN (Bonnarel et al. 2011).

The source code of SOFIA-X has been made publicly available and can be accessed on GitHub.⁶ Because SOFIA-X acts as a wrapper around SOFIA 2, it can be installed independently and is not required for running just a single instance of SOFIA 2.

4 ALGORITHMS

SOFIA 2 largely uses the same algorithms as its predecessor. For that reason we will refrain from a detailed description of the individual algorithms and instead refer the reader to Serra et al. (2015) and references therein. Detailed information about the different algorithms and settings is also available from the official SOFIA 2 user manual that can be obtained from the SOFIA 2 GitHub wiki. Here, we will give a brief overview of some of the core modules of SOFIA 2 with a particular focus on algorithms that are relevant to the analysis presented in this paper or have since been added or changed. At the time of writing, a few useful features from SOFIA 1 have not

⁴<http://www.ivoa.net/documents/TAP/>

⁵<http://www.ivoa.net/documents/ADQL/>

⁶<https://github.com/AusSRC/SoFiAX/>

yet been incorporated into SOFIA 2 (see Section 2) and are expected to be implemented in the near future.

4.1 Input and preconditioning

Like its predecessor, SOFIA 2 currently only supports input data in the *Flexible Image Transport System* (FITS) format (Pence et al. 2010). Data files must have 2–4 dimensions as long as the fourth axis has a size of no greater than one (e.g. Stokes I). It is possible to read and process only a subregion of a data cube. In addition to the actual data cube to be searched, the user can also specify an optional noise cube, weights cube, or gain cube. These will be applied to the data cube prior to either source finding or source parametrization, as appropriate.

Users can manually specify a set of rectangular regions to be flagged prior to source finding. In addition, a new autoflagging algorithm is available which will dynamically flag spatial pixels or spectral channels for which the noise deviates from the median noise across all pixels or channels by a user-specified multiple of the rms as estimated from the more robust median absolute deviation. The autoflagger is intended for the automatic flagging of corrupted data, including channels affected by radio–frequency interference or pixels containing residual continuum emission.

If for some reason the noise level across the input data cube is not constant, then SOFIA 2 will first need to normalize the noise level before being able to apply a constant threshold to the data for the purpose of source finding. This can be achieved by providing either a noise cube by which the data cube will be divided, or a weights cube (holding the inverse of the variance) by the square root of which the data cube will be multiplied. If neither of those are available, or there are residual noise variations not accounted for by the noise or weights cube, then SOFIA 2 can offer to measure and divide by the local noise level in a running window.

4.2 Source finding

SOFIA 2 offers two different source finding algorithms: a simple threshold finder and the *smooth and clip* (S + C) finder. The threshold finder will simply apply an absolute or relative (to the noise) flux threshold to the data and is rarely useful unless the data have already been preconditioned before being read into SOFIA 2. The most powerful and default source finding algorithm of SOFIA 2 is the S + C finder which is described in detail in Serra et al. (2012b). It essentially works by iteratively smoothing the data cube on multiple spatial and spectral scales to extract statistically significant emission above a user-specified detection threshold on each scale. The output from both source finding algorithms will be a binary mask of detected pixels.

An important feature of SOFIA 2 is its ability to pick up both positive and negative flux density values in excess of the source finding threshold. This strategy avoids creating a positive flux bias that would inevitably arise if only positive flux density values were added to the source mask. In addition, the resulting false detections with negative total flux can be used to estimate the statistical reliability of detections with positive flux (see Section 4.4).

4.3 Linking

The purpose of the linker is to combine the detected pixels in the binary mask from the source finder into individual, coherent detections. For this purpose, the linker uses a simple friends-of-friends algorithm that links all pixels within a user-specified merging

radius and assigns a unique identifier to each detection. This linking of pixels occurs both in the spatial plane and along the spectral axis, i.e. sources are effectively treated as three-dimensional collections of pixels. In addition, user-specified minimum and maximum size filters can be applied in the spatial and spectral dimension to remove very large or very small detections that are likely to be false detections due to noise peaks or large-scale artefacts in the data.

4.4 Reliability measurement

One of the most important features of SOFIA 2 is its ability to statistically determine the reliability of each detection and use this information to automatically remove detections from the source catalogue that are deemed unreliable (e.g. Dickinson et al. 2004; Yan & Windhorst 2004; Kovač, Oosterloo & van der Hulst 2009). The underlying algorithm is described in detail in Serra et al. (2012a) and works by comparing the number density of detections with negative and positive total flux in a 3D parameter space made up of the peak flux density, the summed flux density, and the mean flux density across the source.

The reliability measurement is particularly powerful as it allows false positives to be automatically removed from the source catalogue by simply applying a reliability threshold. This allows SOFIA 2 to be pushed deeper than other source finders by applying extreme detection thresholds of as low as three times the noise level without being overwhelmed by the large number of false positives that would otherwise result from such a choice.

In addition to the reliability threshold, the user can also set a signal-to-noise threshold. All detections with an integrated signal-to-noise ratio (SNR) below that threshold will be discarded as unreliable irrespective of their measured reliability. This feature relies on accurate beam information to be present in the data cube header and can be used to automatically remove detections that are too faint to be considered reliable and were assigned a high reliability value purely because of low number statistics in poorly populated regions of parameter space.

Users are advised that the integrated SNR of sources in the final source catalogue could potentially be lower than the threshold defined in the reliability module as the reliability calculation is carried out on the noise-normalized data cube rather than the original cube. In addition, optional features such as mask dilation have the potential to alter the source mask and hence the measured SNR of all detections.

4.5 Parametrization and output

SOFIA 2 will provide measurements of basic source parameters across the three-dimensional source mask. By default, these will be specified in the native pixel coordinates and flux units of the data cube, although the user has the option of manually enabling certain physical corrections, including conversion from pixel coordinates (x , y , and z) to proper world coordinates (e.g. celestial coordinates and frequency in the native units of the data cube) and division of spatially integrated flux parameters by the beam solid angle. SOFIA 2 will extract the required data cube axis descriptors and beam information from the FITS file header, and it is the user's responsibility to ensure that the header information is correct and adequate. In addition, SOFIA 2 implicitly assumes that the spectral channels of the data cube are uncorrelated. If this were not the case, e.g. due to spectral smoothing, then the user would have to manually correct the relevant source parameters to account for the correlation of spectral channels.

SOFIA 2 will also derive basic statistical uncertainties for several fundamental parameters, including the centroid and integrated flux




Figure 2. Left: Resident memory usage of SOFIA 2 relative to the input data volume of 724 MB as a function of elapsed time for a single-threaded run (teal) and a multithreaded run using eight threads (magenta) of SOFIA 2 in comparison to a similar run of SOFIA 1 (grey) on a laptop computer equipped with Intel Core i7–2720QM CPUs clocked at 2.2 GHz. Right: Same but for a run of SOFIA 2 on a 12.6 GB data cube using 16 threads (magenta) versus 1 thread (teal) on one of the ICRAR Hyades cluster nodes equipped with Intel Xeon Gold 6126 CPUs clocked at 2.6 GHz.

of the source. It should be emphasized that these are calculated assuming Gaussian error propagation and may not be representative of the true uncertainties which are often dominated by systematic rather than stochastic errors.

SOFIA 2 can save the final source catalogue in three different formats: plain-text format intended for basic visual inspection, VO-compliant XML format for use in Virtual Observatory tools, and SQL format for integration into a data base. SOFIA 2 can also provide more advanced data products for scientific and diagnostic purposes, including a copy of the source mask, global moment maps of all detections, diagnostic plots created by the reliability module, and data products such as moment maps and integrated spectra for each individual detection.

5 SPEED AND MEMORY USAGE

Two of the most fundamental metrics of a source finding pipeline are the time it needs to process a given amount of data and the peak memory usage relative to the size of the data cube. We determined both metrics on a standard laptop computer with 8 GB of RAM and four Intel Core i7–2720QM CPUs clocked at 2.20 GHz with a total of 8 threads. The test was carried out on a data cube from Serra et al. (2012a) containing genuine noise from an HI observation taken with the Westerbork Synthesis Radio Telescope and several injected galaxies obtained from the WHISP survey (Swaters et al. 2002). The data cube has a size of 360×360 spatial pixels and 1464 spectral channels, corresponding to a data volume of 724 MB. Two identical runs of SOFIA 2 were carried out: one multithreaded run using all eight threads, and another single-threaded run with multithreading disabled altogether. The test was carried out using version 2.2.1 of the software. In addition, we ran a comparable test on the same data cube using SOFIA 1.3.3 to demonstrate the substantial improvement in speed and memory usage of SOFIA 2.

As the time required to read the data cube from disc into memory will depend on several external factors, the test was run with the data already cached in memory in order to measure just the processing time without the I/O contribution. We enabled local noise scaling across a spatial and spectral window size of 31 pixels/channels. We then employed the S + C finder using three spatial filters of 0, 5, and 10 pixels and four spectral filters of 0, 3, 7, and 15 channels.

The flux detection threshold was set to 3.5 times the noise level in each smoothing iteration. In addition, we ran the linker, reliability filter, and parametrizer before writing out the resulting catalogue to disc in all three output formats supported by SOFIA 2. In addition, we created global moment maps as well as standard output products for each individual source (including subcubes, moment maps, and spectra).

The results of the performance test are presented in the left-hand panel of Fig. 2. The resident memory usage of SOFIA 2 peaks at just over 2.25 times the input data volume of 724 MB. This is exactly as expected for a 32-bit data cube, as the pipeline will need to hold at most two copies of the input cube plus an 8-bit source detection mask in memory at any one time. It also demonstrates the significant reduction in memory usage compared with SOFIA 1, which requires memory equivalent to more than five times the input cube size.

It took SOFIA 2 about 136 s to complete the single-threaded run, corresponding to a processing rate of 19 GB per hour. With multithreading enabled, the processing time reduces to just 39 s, corresponding to 65 GB per hour. Multithreading therefore significantly reduces the overall processing time by a factor of 3.5 on a modest system with just eight threads and despite the fact that not all of the algorithms in SOFIA 2 can be, and have been, multithreaded.

In comparison, SOFIA 1 required about 201 s to complete the test, which corresponds to 12.7 GB per hour and indicates that SOFIA 2 is more than five times as fast as SOFIA 1 even on a modest laptop computer and at the same time requires less than half the amount of memory.

It should be noted that these performance measures were obtained with the input data cube already cached in memory, and additional time would be required to read the data from disc. The obtainable I/O speed will vary significantly as a function of several parameters, including the type of storage hardware used, the data access pattern (full cube versus subregion), and any simultaneous I/O activity by other processes in the background. Hence, the performance measures obtained in our test are not expected to scale linearly with data cube size, and larger cubes may require disproportionately more time to process.

To assess the scalability of our performance test results, we repeated the speed and memory test with the same settings on a much larger data cube of 12.6 GB, this time using one of the nodes

Table 1. Run time of SOFIA 2 on a 12.6 GB test data cube with multithreading enabled. The columns show the number of parallel threads used (n_{thr}), the total run time (t_{run}) excluding the time required to read the data cube, the total CPU time (t_{cpu}), the speed-up factor relative to the run with two threads (η), and the resulting parallel fraction of the code (f_p) using equation (1).

n_{thr}	t_{run} (s)	t_{cpu} (s)	η	f_p
2	1925	3291	1.00	–
4	1155	3210	1.67	0.80
8	772	3275	2.49	0.80
16	582	3425	3.31	0.80
32	474	3720	4.06	0.80

of the Hyades computing cluster at the International Centre for Radio Astronomy Research (ICRAR) in Perth, Western Australia, which is equipped with Intel Xeon Gold 6126 CPUs clocked at 2.6 GHz. The results are presented in the right-hand panel of Fig. 2. When using just a single thread, SOFIA 2 completed the run in about 39 min, corresponding to a data rate of 19 GB per hour. When utilizing 16 threads simultaneously, the processing time decreased by a factor of 3.9 to about 10 min, implying a data rate of approximately 76 GB per hour. These values are in good agreement with those from our earlier test on a laptop computer, with a small improvement from the increase in the number of threads from 8 to 16. This suggests that the processing speed of SOFIA 2 remains at a comparable level when scaling up the data volume by a factor of almost 20, and our performance measurements should therefore be robust and scalable.

We can also use our run time measurements to estimate the parallel fraction of the code. In order to achieve this, we repeated the test run on the 12.6 GB data cube while varying the number of parallel threads available to SOFIA 2. The resulting measurements are shown in Table 1. As before, we exclude the initial time required to read in the data cube from the measured run time, t_{run} . We then apply Amdahl’s law (Amdahl 1967),

$$\frac{1}{\eta} = 1 - f_p \left(1 - \frac{1}{\Delta n_{\text{thr}}} \right), \quad (1)$$

to measure the parallel fraction, f_p , of the code, where η is the speed-up and Δn_{thr} is the factor by which the number of threads has increased. With $f_p \approx 0.8$, our measurements indicate that approximately 80 per cent of the code is running in parallel, which is an excellent outcome. This result also suggests that SOFIA 2 should ideally utilize between 8 and 16 threads, beyond which any additional gains in speed become insignificant.

Lastly, users should bear in mind that processing times, memory requirements, and parallelization efficiency will critically depend on the actual algorithms and parameter settings chosen by the user, and specific source finding runs could require more time and memory than the test run presented here. However, we deliberately activated settings that would be typical for a WALLABY source finding run, and the test result should therefore be fairly representative for the typical needs of extragalactic H I source finding.

6 COMPLETENESS, RELIABILITY, AND PARAMETRIZATION ACCURACY

In order to assess the performance of SOFIA 2 with respect to completeness, reliability, and parametrization accuracy, we created a mock data cube by injecting model galaxies into a data cube

containing genuine ASKAP noise. The noise was obtained by extracting a subregion of 1501×1501 spatial pixels and 1501 spectral channels from a WALLABY pre-pilot data cube of the Eridanus galaxy cluster, resulting in a file size of about 12.6 GB. The spatial pixel size of the cube is 6 arcsec (with a beam size of ≈ 30 arcsec), and the spectral channel width is about 18.5 kHz which corresponds to a velocity resolution of about 4 km s^{-1} at redshift 0. The noise cube was extracted from the frequency range of 1323–1351 MHz to minimize the risk of contamination with genuine H I emission. We then used the noise scaling algorithm of SOFIA 2 to normalize the noise level to 1, thus allowing us to conveniently specify all source parameters in units of SNR.

Next, we generated several thousand model galaxies using the GALMOD task of the GIPSY data processing software (Allen et al. 2011). To ensure that the resulting galaxies would cover a wide range of different observational properties, we randomly varied several galaxy parameters within a meaningful range, including peak H I column density (10^{20} – 10^{21} cm^{-2}), rotation velocity (30 – 220 km s^{-1}), exponential scale length on the sky (4.5 – 36 arcsec), disc inclination (0° – 85°), and position angle on the sky (0° – 360°).

The model galaxies were convolved with a 30 arcsec Gaussian beam in accordance with the expected restoring beam of the WALLABY data. 3200 model galaxies were then injected into the noise cube on a regular grid in an attempt to fit as many model galaxies as possible. The flux density of each model galaxy was scaled by a constant factor to ensure that the integrated SNRs of most of the galaxies would fall into the range of about 0–10, as this is the most interesting range within which we would expect the completeness to increase from 0 to 1 (at $\text{SNR} \approx 5$).

In addition, we created a second version of the data cube by using the same set of galaxies, but this time using five times the original flux scaling factor to extend our sample into the higher SNR range of up to about 50 for the purpose of checking the source parametrization accuracy of SOFIA 2. The two SNR samples combined therefore provide us with an overall sample of 6400 mock galaxies.

Both model data cubes were then processed with SOFIA 2 using the S + C finder with spatial kernels of 0, 5, and 10 pixels and spectral kernels of 0, 3, 7, and 15 channels. In addition, we set the detection threshold to 3.5 times the noise level, the linker radius to 2 pixels/channels, and the minimum source size to eight spatial pixels and five spectral channels. Lastly, we enabled the reliability filter to automatically remove unreliable detections, setting a reliability threshold of 0.9, a kernel scale factor of 0.4, and a signal-to-noise threshold of 2.8.

6.1 Completeness and reliability

Among the most fundamental performance indicators of any source finding algorithm are its *completeness* and *reliability*. Completeness, C , is defined as the fraction of genuine sources being successfully detected by the source finder, hence

$$C = \frac{N_{\text{gen}}}{N_{\text{tot}}}, \quad (2)$$

where N_{gen} is the number of genuine sources detected, while N_{tot} is the total number of sources present. Likewise, reliability, R , is defined as the fraction of detected sources that are genuine, N_{gen} , as opposed to false positives due to artefacts or noise, N_{false} , hence

$$R = \frac{N_{\text{gen}}}{N_{\text{gen}} + N_{\text{false}}}. \quad (3)$$

Completeness and reliability are most meaningfully defined in differential form as a function of integrated SNR, with both completeness and reliability expected to approach 100 per cent at high SNR, while gradually decreasing towards low SNR.

Likewise, completeness and reliability of a source finding run will strongly vary with the parameter settings of the source finder, most notably the detection threshold. Using a lower detection threshold will increase the number of genuine detections at low SNR and hence completeness, while at the same time producing more false detections and thus decreasing reliability. Therefore, the main challenge with automated source finding is to establish optimal settings that strike a balance between acceptable levels of completeness and reliability.

6.1.1 Reliability

In order to obtain an estimate of the expected reliability of SOFIA 2, we first ran the pipeline on the original noise cube *without* injected mock galaxies, using exactly the same settings as for the two data cubes with galaxies injected. If SOFIA 2 were fully reliable, this experiment should yield no detections as the cube should only contain stochastic noise, albeit genuine noise produced by the ASKAP telescope and receiver system.

Overall, SOFIA 2 reported 12 328 detections at the 3.5σ detection threshold chosen for this experiment. Of these, 6370 have positive total flux, while 5958 have negative flux. Virtually all of these detections are deemed unreliable by the reliability filter in SOFIA 2, and only a single detection with positive flux remains after reliability filtering. Within just 5 arcsec of the measured sky position of that detection of $\alpha = 03^{\text{h}}38^{\text{m}}34^{\text{s}}$ and $\delta = -22^{\circ}46'08''$ there is a bright optical counterpart, LEDA 809162, which is classified as a galaxy in the NASA/IPAC Extragalactic Database.⁷ This suggests that the sole signal detected by SOFIA 2 in the noise cube is a genuine high-redshift H I emission line at $z = 0.066$ associated with LEDA 809162. Unfortunately, no optical redshift measurement is available for LEDA 809162 to unambiguously confirm that the H I signal is genuine and indeed associated with that galaxy.

With the sole detection being almost certainly a genuine galaxy at higher redshift, we can therefore be confident that no false positives were picked up by SOFIA 2 after reliability filtering, indicating that the output catalogue produced by SOFIA 2 is 100 per cent reliable at all SNR levels. This outcome highlights the excellent quality of ASKAP data as well as the power of the reliability filter implemented in SOFIA 2 which appears to be capable of accurately discarding even a large number of false positives generated as a result of the low detection threshold of 3.5σ applied in this experiment.

It should be emphasized at this point that the reliability filter in SOFIA 2 is based on the assumption that the image noise is symmetric about zero. The reliability of the source catalogue from SOFIA 2 will therefore critically depend on how clean the underlying data are. In particular, effects such as interference, residual continuum emission, H I absorption, etc. are likely to reduce the effectiveness of the reliability filter, potentially reducing both the reliability and completeness of the resulting source catalogue. Later on in Section 7, we will see an example of an actual ASKAP data set that is not as clean as the ASKAP noise used in our mock data set, resulting in a reduced effectiveness of the reliability filter and hence a somewhat lower reliability of the source catalogue.

⁷<https://ned.ipac.caltech.edu/>

6.1.2 Completeness

In order to analyse the completeness of the source finding run as a function of SNR, we will first need to establish a way of characterizing the integrated SNR of a source. The integrated flux of a source is given by

$$F = \frac{\Delta\nu}{\Omega} \sum_{i=1}^N S_i, \quad (4)$$

where $\Delta\nu$ is the width of a frequency channel, Ω is the solid angle of the restoring beam in units of pixels, and S_i is the flux density value of pixel, i , with the summation carried out over all N pixels considered to be part of the source. For a Gaussian beam, $4 \ln(2) \Omega = \pi a b$, with a and b being the full width at half-maximum of the major and minor axis of the beam in units of pixels. Note that we have made the explicit assumption that the spectral channels of the data cube are uncorrelated, which is certainly the case with our ASKAP data.

In principle, we can use the Gaussian error propagation law to determine the statistical uncertainty of the flux measurement, thus

$$\sigma_F = \frac{\sqrt{N} \Delta\nu \sigma_{\text{rms}}}{\sqrt{\Omega}}, \quad (5)$$

where σ_{rms} is the original noise level per pixel and we have additionally accounted for the fact that spatial pixels are partially correlated due to the finite size of the beam, while we assume spectral channels to be entirely uncorrelated. The integrated SNR is then simply given as $\text{SNR} = F/\sigma_F$.

While this method works well with detected sources, yielding the observed SNR, SNR_{obs} , it will be more challenging to apply to non-detections or mock galaxies. A practical solution would be to sum the pixels of the mock source in the order of decreasing flux density until a certain fraction of the total flux is reached. Using a cut-off of 99 per cent of the total flux, we have made use of this method here to define SNR_{99} for the purpose of calculating completeness. A more detailed description of the definition of SNR_{99} and its relation to the observed SNR is given in Appendix A.

The completeness resulting from our mock galaxy experiment is shown in Fig. 3 as a function of SNR_{99} . Apparently, SOFIA 2 is capable of achieving 90 per cent completeness at $\text{SNR}_{99} \approx 5$, with all galaxies above $\text{SNR}_{99} = 6$ being picked up. Interestingly, SOFIA 2 will still detect a few galaxies at very low SNRs in the range of 2–3 that must have been boosted by contributions from collocated noise peaks. Even at $\text{SNR}_{99} = 3.5$ the completeness is still sitting at the 50 per cent level.

The outcome of the mock galaxy experiment demonstrates that the S + C source finding algorithm of SOFIA 2 in combination with its sophisticated reliability filter is capable of extracting a highly complete and reliable source catalogue from a large H I data cube. As said before, the performance of SOFIA 2 will stand and fall with the ability of the reliability filter to remove false detections. The cleaner the data, the better the algorithm is expected to perform, while the presence of artefacts will likely degrade the reliability and completeness of the source catalogue.

Nevertheless, the experiment presented here was carried out using real noise from an actual spectral-line observation with ASKAP demonstrating that optimal performance can be achieved on real data and that SOFIA 2 is in principle ready to handle H I data from the WALLABY survey.




Figure 3. Stacked histogram of detected (yellow) and undetected (grey) mock galaxies and the resulting differential completeness (black curve) as a function of SNR_{99} , demonstrating that SOFIA 2 is able to achieve 90 per cent completeness at $\text{SNR}_{99} \approx 5$.




Figure 4. Position measurement errors colour coded by SNR_{99} of each mock galaxy. The black, dotted circle marks the size of the 30-arcsec beam.

6.2 Parametrization accuracy

Another important performance metric to assess is the accuracy with which SOFIA 2 is able to recover basic observational parameters of the mock galaxies, such as position, integrated flux, or spectral profile width. As with completeness and reliability, this is expected to be a function of integrated SNR, as measurements of fainter sources will suffer from larger statistical uncertainties and potentially be more susceptible to systematic errors such as the loss of emission from the faint outer regions of a source.

6.2.1 Source location

The position measurement errors from the mock galaxy experiment are presented in Fig. 4. As can be seen, the positions produced by SOFIA 2 are on average highly accurate with a mean of 0.0 ± 1.6 pixels

in both x and y , the standard deviation corresponding to about one-third of the size of the beam. As expected, the position accuracy is a function of integrated SNR of the source, with a standard deviation of about 0.5 pixels for brighter galaxies of $\text{SNR}_{99} > 10$ as compared to about 2.1 pixels for fainter galaxies of $\text{SNR}_{99} < 10$.

While most mock galaxies are located within the size of the beam, Fig. 4 does reveal a faint halo of points beyond the beam size. This halo is likely to have been caused by sources that were either just partially detected or fragmented into multiple components by SOFIA 2. This is known to occasionally occur in the case of very faint, edge-on galaxies where the two halves of the galaxy are picked up as separate detections.

The situation is very similar for the location accuracy along the frequency axis, where the mean error is 0.0 ± 8.6 channels. Again, the standard deviation is affected by outliers due to partially detected or fragmented galaxies at low SNR. When only considering galaxies with $\text{SNR}_{99} > 10$, the standard deviation drastically reduces to just 0.9 channels demonstrating the excellent recovery of frequency centroids by SOFIA 2.

6.2.2 Fragmented sources

To further investigate the issue of fragmented detections, we plot in Fig. 5 the ratio of spectral position error and spectral line width, r_z/w_{20} , against the ratio of spatial position error and spatial major axis size, r_{xy}/w_{maj} , for all mock galaxies detected by SOFIA 2. There is a strong concentration of sources near the origin that are likely to have been left intact by the source finder as their position errors are very small compared to their spatial and spectral extent.

In contrast, the population of fragmented sources is clearly visible as the extended halo of detections occupying the positional error range of roughly 20–50 per cent of the source extent. In order to separate between intact and fragmented sources, we apply a conservative threshold in the form of an elliptical radius of $r_{xy} < 0.25 w_{\text{maj}}$ and $r_z < 0.125 w_{20}$, as marked by the dashed curve in Fig. 5. We consider all detections within that threshold to be intact (grey data points), while most, if not all, of the fragmented sources are expected to be located outside that threshold (teal data points). We will maintain this cut and the associated colour scheme throughout the remainder




Figure 5. Ratio of spectral position error and spectral line width, r_z/w_{20} , plotted against the ratio of spatial position error and spatial major axis size, r_{xy}/w_{maj} . The black, dashed curve marks the elliptical boundary of $r_{xy} < 0.25 w_{maj}$ and $r_z < 0.125 w_{20}$ that we use to separate sources that have been left intact (grey data points) from sources that are likely to have been fragmented (teal data points) by SOFIA 2.

of this section to be able to separate between intact and fragmented sources in our analysis.

The fraction of sources that have potentially been fragmented by SOFIA 2 is a strong function of SNR. At $\text{SNR}_{99} = 5$ approximately 34 per cent of the detections are located outside the relative position error threshold considered here, whereas that fraction effectively drops to zero beyond $\text{SNR}_{99} = 10$, suggesting that sources near the detection threshold are most susceptible to being fragmented, as would be expected.

Likewise, the fraction of fragmented sources varies strongly with spectral line width. While less than 10 per cent of galaxies with narrow lines of $w_{20} \lesssim 150 \text{ km s}^{-1}$ are getting fragmented in our specific experiment, that fraction gradually increases to about 40 per cent for galaxies with a broad spectral profile of $w_{20} \gtrsim 400 \text{ km s}^{-1}$. Again, this is expected at low SNR levels, as edge-on galaxies with large rotation velocities are most susceptible to being fragmented by the source finder due to their broad double-horn profiles.

As fragmented sources have the potential to significantly alter the results of certain scientific studies, e.g. measurements of the H I mass function, they would need to be handled separately. Identification of partially detected or broken-up sources could potentially be achieved through comparison with optical catalogues or identification of close pairs of detections with similar properties.

6.2.3 Peak flux density

The peak flux density of a source reported by SOFIA 2 is simply the highest flux density value encountered within the 3D source mask (not the peak flux density of the integrated spectrum). As such, the peak flux density will be strongly affected by the presence of noise which should result in a significant positive bias. This is illustrated in the upper left panel of Figs 6 and 7 where the ratio of measured versus true peak flux density of the mock galaxies is plotted as a

function of SNR_{99} and SNR_{obs} , respectively.⁸ The significant bias, in particular below $\text{SNR}_{99} \approx 10$, means that the peak flux density is not a particularly meaningful parameter beyond the purpose of basic sanity checking, and its use in any form of scientific analysis is therefore discouraged.

6.2.4 Integrated flux

One of the most fundamental observational parameters to be extracted from any source finding effort is the integrated flux of a source as it is required for determining the H I mass of a galaxy. SOFIA 2 measures the integrated flux, F , by summing the flux density values, S_i , of all N pixels contained in the source mask, additionally multiplying by the spectral channel width, $\Delta\nu$, and dividing by the beam solid angle, Ω , if explicitly requested to do so by the user (equation 4).

In addition, SOFIA 2 calculates the statistical uncertainty of the integrated flux measurement by assuming Gaussian error propagation and correcting for the fact that spatial pixels will be correlated due to the finite beam size (equation 5). SOFIA 2 measures the local noise level in the vicinity of each detection for this purpose.

The flux measurement errors resulting from the mock galaxy experiment are presented in the upper right panel of Figs 6 and 7 as a function of SNR_{99} and SNR_{obs} , respectively. For galaxies not broken up into multiple detections (grey and black data points), the integrated flux is accurately recovered across the entire range of SNRs, with the statistical uncertainty increasing towards lower SNR_{99} , as expected. A closer inspection of Fig. 6 reveals that the mean relative flux measurement error is not exactly centred on zero. Instead, the flux measurement from SOFIA 2 is on average about 2–3 per cent higher across most of the SNR range. This effect is particularly noticeable at higher SNR_{99} where the statistical uncertainties are smaller.

Our investigation of this flux measurement bias revealed that it is likely caused by a particular aspect of the S + C finder whereby at the beginning of each iteration all pixels already detected in previous iterations are set to $\pm n$ times the noise level (preserving the original sign of the flux value) in the data cube before smoothing to prevent the smoothing operation from smearing out the emission beyond the boundaries of the source. A side effect of this approach appears to be that the resulting source mask tends to be grown in directions where a predominantly positive contribution from the noise can be expected, thus resulting in a net positive flux bias.

The effects appear to depend directly on the replacement value chosen. The default value of $n = 2$ used here will result in the observed 2–3 per cent positive flux bias, while we have found the bias to disappear entirely when slightly reducing the value to 1.5, albeit at the cost of a somewhat stronger negative flux bias at the faint end of the SNR range. If high flux accuracy of better than 3 per cent is required, a statistical correction would need to be applied to remove any bias.

In Fig. 8, we show a histogram of the flux measurement error divided by the flux measurement uncertainty from SOFIA 2. Under perfect conditions, err_F/σ_F should have a Gaussian distribution with a centroid of 0 and a standard deviation of 1. A Gaussian fit to the histogram yields a centroid of 0.27 ± 0.02 and a standard deviation of 1.22 ± 0.02 , which is very close to the expected values and indicates that the flux uncertainties reported by SOFIA 2 are accurate. The small offset and slight broadening of the distribution are not unexpected and likely reflect the minor flux bias and its variation with SNR.

⁸Parametrization accuracy is presented here as the ratio of measured versus true value, which is equal to the relative error plus one and better accounts for the multiplicative nature of the flux measurement process in radio astronomy.




Figure 6. Ratio of measured versus true source parameter as a function of SNR_{99} for peak flux density (upper left), integrated flux (upper right), w_{20} line width (lower left), and w_{50} line width (lower right). Intact sources are shown in grey, fragmented ones in teal. The black data points and error bars show the mean and standard deviation of the grey data points in intervals of 2.5.

6.2.5 Spectral line width

Another fundamental observational parameter to be extracted by any source finder is the width of the integrated spectral profile of a galaxy which forms the basis for determining the rotation velocity. SOFIA 2 determines the integrated spectral profile by summing over the spatial pixels contained in the source mask in each spectral channel. The w_{20} and w_{50} line widths are then measured by moving inwards from both ends of the spectrum until the signal exceeds 20 per cent or 50 per cent of the peak flux density in the spectrum, respectively. For improved accuracy, SOFIA 2 will linearly interpolate across the two channels in between which the signal exceeds the threshold for the first time.

As the line width measurement depends on the peak flux density of the spectrum, it is systematically affected by the presence of noise which will result in an overestimation of the peak flux density and thus an underestimation of the line width. This particularly affects the w_{50} measurement and is one of several reasons for why w_{50} may not provide an accurate measurement of the line width and rotation velocity of a galaxy at low SNR (see also Tully & Fouqué 1985; Ho 2007). The measured w_{50} errors of the mock galaxies are plotted in the lower right panel of Figs 6 and 7 as a function of SNR_{99} and SNR_{obs} , respectively. The significant negative bias, particularly at $\text{SNR}_{99} \lesssim 10$, is clearly visible, and we therefore discourage the use of w_{50} when dealing with objects of low SNR.

The w_{20} measurement errors resulting from the mock galaxy experiment are presented in the lower left panel of Figs 6 and 7

as a function of SNR_{99} and SNR_{obs} , respectively. Unlike w_{50} , the w_{20} measurement from SOFIA 2 is highly accurate across the entire range down to $\text{SNR}_{99} \approx 5$, again with the exception of those galaxies that got fragmented by the source finder. The measurement of w_{20} is even more accurate than the integrated flux measurement, and there is no sign of any significant bias. As expected, the statistical uncertainty increases towards the faint end.

One of the main reasons for why the measurement of w_{20} by SOFIA 2 is so accurate is the fact that the integrated spectrum is created across the three-dimensional source mask, the shape of which can be expected to closely match the morphology of the source. This has the effect of maximizing the SNR and suppressing noise near the edges of the spectrum, where only few spatial pixels are expected to contribute. As a consequence, the w_{20} measurement is not quite as much affected by noise as in classical studies where the spectrum is usually integrated across a rectangular mask and thus much noisier near the edges of the profile.

7 WALLABY PILOT DATA

To demonstrate the fitness of SOFIA 2 for full-scale WALLABY source finding, we ran a single-threaded version of the pipeline on the first WALLABY pilot survey data in the direction of the Hydra galaxy cluster using the parallel SOFIA-X framework introduced in Section 3. Approximately 85 GB of data in a field of about $3.0^\circ \times 5.5^\circ$




Figure 7. Same as Fig. 6 but as a function of observed SNR, SNR_{obs} .




Figure 8. Histogram of the true integrated flux error, $\text{err}_F = F_{\text{obs}} - F_{\text{model}}$, divided by the integrated flux uncertainty, σ_F , derived by SOFIA 2 for the mock galaxies. The dashed, red line shows the result of a Gaussian fit to the histogram which should ideally be centred at 0 with a standard deviation of 1. The vertical, dashed, grey line marks the actual peak position of 0.27, which is very close to the expected value.

(equivalent to 1801×3601 pixels) on the sky out to a redshift of $z \approx 0.05$ (equivalent to 3517 frequency channels) were processed. The task was split across four nodes of the Hyades computing cluster at ICRAR. The output catalogues and images were directly written into

a dedicated source data base, and duplicate detections in the 300-pixel overlap bands between the four subregions were automatically identified and merged.

Overall, SOFIA 2 completed the source finding run in approximately 2 h (without multithreading), producing 197 unique detections after the automatic removal of duplicates. As a result of artefacts in the image data produced by the ASKAP data reduction pipeline (Guzman et al. 2019), some of these detections were found to be unlikely to be genuine astronomical sources and manually removed from the catalogue and output images. 148 detections remained after this process, and the moment-0 image resulting from the test run is presented in panel A of Fig. 9. The faintest sources in this image have an *observed* integrated SNR of about 4. These sources are so faint that they are actually no longer visible to the naked eye in an individual channel map of the data cube as their signal remains well below the noise level.

To illustrate this effect, we show in panel B of Fig. 9 an individual channel map from the original data cube in the direction of the galaxy NGC 3313. The source mask produced by SOFIA 2 in this frequency channel (bold, black contour) extends well beyond the brightest knots in the H I disc of NGC 3313. The fainter regions of the disc are not visible at the original resolution of the data as they are sitting below the noise level. Yet, thanks to its spatial and spectral smoothing, the S + C algorithm implemented in SOFIA 2 is capable of detecting the full extent of the disc within the limits set by the sensitivity of the data. This is illustrated by the white contour in panel B which corresponds to approximately three times the noise level of the data cube after spatial and spectral smoothing with the largest kernels that




Figure 9. Panel A: Moment-0 map from a SOFIA 2 run on a $3.0^\circ \times 5.5^\circ$ section of the WALLABY pilot survey field in the direction of the Hydra galaxy cluster. Almost 150 individual galaxies were detected by SOFIA 2. Note that only pixels contained within the source mask are included in the moment calculation, and all other pixels are set to zero. Panel B: An individual channel map from the original data cube in the direction of the galaxy NGC 3313 (see black box in panel A). The bold, black contour shows the outline of the source mask produced by SOFIA 2. The orange and grey contours correspond to ± 6 mJy ($\approx \pm 3\sigma$) in the original data cube. The bold, white contour corresponds to 2 mJy ($\approx 3\sigma$) after smoothing spatially with a Gaussian kernel of 10 pixels FWHM and spectrally with a boxcar filter of 15 channels to mimic the maximum smoothing scales applied by SOFIA 2. Panel C: The resulting moment 0 map of NGC 3313 integrated across the full 3D source mask. The SNR per pixel is shown as contours at levels of -1 (light grey), 1 (dark grey), 3 (bold, red), 5 , and 7 (both black). Note that the additional white contour near the right-hand edge of panel B is a satellite galaxy of NGC 3313 that was also detected by SOFIA 2 but is not included in panel C, as it is a separate source with its own mask. Also see fig. 14 in Koribalski et al. (2020) for a more detailed view of this data cube.

were actually applied by the S + C finder. The SOFIA 2 source mask nicely traces the full extent of the disc after smoothing, as is also shown by the resulting moment 0 map of NGC 3313 across the full source mask in panel C.

The outcome demonstrates that SOFIA 2 is capable of processing real WALLABY data in a short amount of time. In fact, with multithreading enabled, the same run would have been completed in under 30 min if 16 CPU threads had been utilized simultaneously. Such processing times are negligible compared to the 16 h it took to acquire the data at the telescope, let alone the additional substantial amount of time that was needed to calibrate and image the raw visibility data. The false positives picked up by SOFIA 2 are the result of artefacts caused by insufficient flagging of the visibility data and inadequate continuum subtraction by the ASKAP data reduction pipeline. Efforts to improve the flagging and continuum subtraction algorithms are underway and expected to yield much

cleaner data cubes that should get us significantly closer to the 100 percent reliability achieved in our mock galaxy experiment. The mock data were based on real ASKAP noise from WALLABY pre-pilot observations of the Eridanus galaxy cluster which were less affected by such artefacts.

We would also like to emphasize that for the Hydra data we deliberately pushed SOFIA 2 to its limits by adopting a significantly lower reliability threshold (0.6 versus 0.9) and smaller reliability kernel scale factor (0.3 versus 0.4) than for the mock data test presented in Section 6. This is expected to improve the completeness of the catalogue as compared to the mock data experiment but at the expense of a reduction in reliability. The lower reliability obtained for the Hydra data is therefore not in tension with the mock data experiment but rather a deliberate choice in our attempt to push for even greater completeness by challenging the limits of SOFIA's algorithms. If we had adopted the more conservative and robust

settings of the mock data experiment, as will likely be the case for the full WALLABY survey, we would have obtained completeness and reliability values more in line with those presented in Section 6.

Lastly, the test run also demonstrates the ability of the parallel SOFIA-X framework to split large-scale WALLABY data sets across multiple computing nodes and automatically merge the resulting catalogues and data products without any conflict. We are therefore confident at this stage that SOFIA 2 fulfils all of the requirements for full-scale and real-time source finding on the large data volumes expected to be produced by the WALLABY survey in the near future.

8 SUMMARY AND CONCLUSIONS

SOFIA 2 is a new extragalactic HI source finding pipeline that is significantly faster and more memory-efficient compared to its predecessor, SOFIA 1, thanks to the use of the C programming language in combination with multithreading of time-critical algorithms. A substantial reduction in the number of third-party dependencies also means that SOFIA 2 is also much easier to install than its predecessor. In addition, we have developed a parallel framework called SOFIA-X that allows the processing of large data cubes to be split across multiple computing nodes in order to reduce processing times even further.

Our performance tests of SOFIA 2 on both mock data and genuine ASKAP data have demonstrated that the pipeline is capable of processing the large data volumes expected to be produced by the WALLABY survey in much better than real time when run in parallel on multiple computing nodes and with multithreading enabled. In our tests, we have been able to achieve processing rates of the order of 76 GB per hour using 16 threads on a single node of the Hyades cluster at ICRAR. This should in principle enable us to process each 800 GB WALLABY data cube in less than 1 h if a sufficiently large machine with about 20 nodes and 16 threads per node could be utilized. Even if the number of available nodes and threads were lower, we would still be able to process the entire cube in much less than real time (16 h observing time plus calibration and imaging), demonstrating that SOFIA 2 is sufficiently fast at this stage to handle full-scale WALLABY survey data. In addition to this substantial increase in speed, SOFIA 2 also occupies significantly less memory than its predecessor, SOFIA 1, thus allowing much larger volumes of data to be loaded and processed simultaneously on a single node.

Tests on mock galaxies injected into genuine ASKAP HI data have shown that SOFIA 2 is in principle capable of achieving close to 100 percent reliability thanks to its built-in reliability measurement and filtering module, while still maintaining about 90 percent completeness at an integrated signal-to-noise level of 5 and 100 percent completeness above $\text{SNR}_{99} \approx 6$. This result is encouraging and suggests that WALLABY will be able to achieve its anticipated detection rate and science goals. Users should note that their data will need to be relatively clean and free from major artefacts for the reliability filter to be fully effective.

The accuracy of basic observational source parameters derived by SOFIA 2 is generally high, and in particular spatial and spectral centroids, integrated fluxes, and w_{20} line widths are accurately recovered, with just minor biases of no more than a few percent across the full range of SNRs considered in our test. Recovery of the peak flux density and w_{50} line width is more strongly affected by the presence of noise in the data and hence subject to significant biases, in particular at lower SNRs; their use in any kind of scientific analysis is therefore not recommended.

There is a general risk of faint sources of $\text{SNR}_{99} < 10$ getting broken up into multiple detections or being only partially detected by SOFIA 2. This is particularly the case for faint, edge-on galaxies

where the HI emission is concentrated in the two ‘horns’ of the spectrum, with very little emission in between. It should be noted that the fragmentation of sources is a direct consequence of the stochastic nature of the noise in combination with the low SNR of some of the objects picked up by SOFIA 2 and therefore in principle unavoidable.

Mitigation strategies should be adopted to deal with both the issue of source fragmentation as well as parametrization bias, as they would otherwise have the potential to corrupt any scientific analysis based on the source finding output. Possible strategies for identifying fragmented sources could involve the identification of close pairs of detections with otherwise similar parameters, or a comparison with optical position and redshift measurements to identify suspicious offsets. Alternatively, a signal-to-noise cut of about 10 can be used to obtain a clean catalogue that is largely unaffected by source fragmentation.

Likewise, if extremely high parametrization accuracy of better than a few percent is required, a full-scale bias assessment using mock sources will need to be carried out similar to the mock data experiment presented here. This would also help with the establishment of completeness and reliability limits. Another possibility of reducing parametrization biases would be the use of external software and algorithms for the parametrization of sources detected by SOFIA 2. As an example, fitting a Gaussian function or the Busy Function (Westmeier et al. 2014) to the integrated spectrum provided by SOFIA 2 could help in obtaining a more accurate measurement of the w_{50} line width.

As SOFIA 2 is likely to be useful for other large-scale HI surveys and potentially other spectral-line data, including observations of the CO line and other molecular transitions, the software has been made freely available to the entire community under an open-source licence and can be downloaded from GitHub. Feedback and bug reports from users of SOFIA 2, either via e-mail or through the official GitHub issue tracking system, are encouraged and will help us to further improve the software and make it as useful and robust as possible for a wide range of spectral-line data.

Lastly, we would like to emphasize again that all of the performance measures presented in this paper, including processing speed, memory usage, completeness, reliability, parametrization accuracy, and source fragmentation fraction, are specific to the data cubes and parameter settings used here and should not be relied upon. Other data sets are likely to require different settings and will almost certainly produce deviating performance metrics. Potential users of SOFIA 2 are therefore advised to perform their own checks using mock data that more closely reflect the characteristics of their specific observational data.

ACKNOWLEDGEMENTS

We wish to thank the referee for not only providing us with a constructive and helpful report but also taking the time to install and test SOFIA 2. This effort has led to several improvements to both the software and its documentation. The Australian SKA Pathfinder is part of the Australia Telescope National Facility which is managed by CSIRO. 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 Centre. Establishment of ASKAP, the Murchison Radio-astronomy Observatory and the Pawsey Supercomputing Centre are initiatives of the Australian Government, with support from the Government of Western Australia and the Science and Industry Endowment Fund. We acknowledge the Wajarri Yamatji

people as the traditional owners of the Observatory site. The authors acknowledge support from the Astronomy Data And Computing Services (ADACS) Software Support Programme and the Australian SKA Regional Centre (AusSRC) Design Study Programme. Parts of this research were supported by the Australian Research Council Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D), through project number CE170100013. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 679627, project name FORNAX). JMVdH acknowledges support from the European Research Council under the European Union's Seventh Framework Programme (FP/2007–2013)/ERC Grant Agreement no. 291531 (HISStoryNU). This research has made use of the NASA/IPAC Extragalactic Database (NED), which is operated by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

DATA AVAILABILITY

The source code of SOFIA 2 is available from GitHub at <https://github.com/SoFiA-Admin/SoFiA-2/>. The source code of SOFIA-X is available from GitHub at <https://github.com/AusSRC/SoFiAX/>. The WALLABY pre-pilot H1 data of the Eridanus cluster are available from the CSIRO ASKAP Science Data Archive (CASDA) at <http://hdl.handle.net/102.100.100/52548> (DOI: 10.25919/0yc5-f769). The WALLABY pilot H1 data of the Hydra cluster are available from CASDA at <http://hdl.handle.net/102.100.100/319460> (DOI: 10.25919/0v41-1055).

REFERENCES

- Allen R. J., Ekers R. D., Terlouw J. P., Vogelaar M. G. R., 2011, *Astrophysics Source Code Library*, record ascl:1109.018
- Amdahl G. M., 1967, Validity of the Single-Processor Approach to Achieve Large Scale Computing Capabilities. AFIPS Joint Spring Conference Proceedings 30 (Atlantic City, NJ, April 18–20). AFIPS Press, Reston, VA, p. 483
- Blue Bird J. et al., 2020, *MNRAS*, 492, 153
- Bonnarel F. et al., 2000, *A&AS*, 143, 33
- Calabretta M. R., 2011, *Astrophysics Source Code Library*, record ascl:1108.003
- Camilo F., 2018, *Nat. Astron.*, 2, 594
- de Blok W. J. G. et al., 2018, *ApJ*, 865, 26
- Dagum L., Menon R., 1998, *Comput. Sci. Eng.*, 1, 46
- Dewdney P. E., Hall P. J., Schilizzi R. T., Lazio T. J. L. W., 2009, *Proc. IEEE*, 97, 1482
- Dickinson M. et al., 2004, *ApJ*, 600, L99
- Flöer L., Winkel B., 2012, *Publ. Astron. Soc. Aust.*, 29, 244
- For B.-Q. et al., 2019, *MNRAS*, 489, 5723
- Guzman J. et al., 2019, *Astrophysics Source Code Library*, record ascl:1912.003
- Ho L. C., 2007, *ApJ*, 669, 821
- Hotan A. W. et al., 2021, *Publ. Astron. Soc. Aust.*, 38, e009
- Jonas J. L., 2009, *Proc. IEEE*, 97, 1522
- Jurek R. J., 2012, *Publ. Astron. Soc. Aust.*, 29, 251
- Koribalski B. S. et al., 2020, *Ap&SS*, 365, 118
- Kovač K., Oosterloo T. A., van der Hulst J. M., 2009, *MNRAS*, 400, 743
- Pence W. D., Chiappetti L., Page C. G., Shaw R. A., Stobie E., 2010, *A&A*, 524, A42
- Popping A., Jurek R. J., Westmeier T., Serra P., Flöer L., Meyer M., Koribalski B. S., 2012, *Publ. Astron. Soc. Aust.*, 29, 318
- Serra P., Jurek R. J., Flöer L., 2012a, *Publ. Astron. Soc. Aust.*, 29, 296
- Serra P. et al., 2012b, *MNRAS*, 422, 1835
- Serra P. et al., 2015, *MNRAS*, 448, 1922

- Swaters R. A., van Albada T. S., van der Hulst J. M., Sancisi R., 2002, *A&A*, 390, 829
- Taylor M. B., 2005, in Shopbell P., Britton M., Ebert R., eds, *ASP Conf. Ser.*, Vol. 347, *Astronomical Data Analysis Software and Systems XIV*. Astron. Soc. Pac., San Francisco, p. 29
- Tully R. B., Fouqué P., 1985, *ApJS*, 58, 67
- Verheijen M. A. W., Oosterloo T. A., van Cappellen W. A., Bakker L., Ivashina M. V., van der Hulst J. M., 2008, in Minchin R., Momjian E., eds, *AIP Conf. Ser.* Vol. 1035, *The Evolution of Galaxies Through the Neutral Hydrogen Window*. Am. Inst. Phys., New York, p. 265
- Westmeier T., Jurek R., Obreschkow D., Koribalski B. S., Staveley-Smith L., 2014, *MNRAS*, 438, 1176
- Whiting M. T., 2012, *MNRAS*, 421, 3242
- Whiting M. T., Humphreys B., 2012, *Publ. Astron. Soc. Aust.*, 29, 371
- Yan H., Windhorst R. A., 2004, *ApJ*, 612, L93
- Yoo A. B., Jette M. A., Grondona M., 2003, in Feitelson D., Rudolph L., Schwiigelshohn U., eds, *Lecture Notes in Computer Science*, Vol. 2862, *Job Scheduling Strategies for Parallel Processing*, JSSPP 2003. Springer, Berlin, p. 44

APPENDIX A: SIGNAL-TO-NOISE RATIO

For the analysis presented in this work, we have made use of SNR_{99} to characterize the SNR of mock galaxies. For this purpose, we first convolve the source model by the actual telescope beam used in the simulation. We then derive SNR_{99} by summing the flux densities in the individual pixels making up the source in the order of decreasing flux density up to the point where 99 per cent of the total flux has been accumulated. SNR_{99} is then obtained by dividing that summed flux density, S_{99} , by the associated statistical uncertainty as defined in equation (5), assuming the actual noise level per pixel, σ_{rms} , from the final mock data cube. Combining equations (4) and (5) thus yields

$$\text{SNR}_{99} = \frac{S_{99}}{\sqrt{N_{99} \Omega \sigma_{\text{rms}}}}, \quad (\text{A1})$$

where Ω is the beam solid angle in units of pixels and N_{99} is the actual number of pixels that had to be summed to obtain S_{99} .

An important question that arises from the use of SNR_{99} is how it relates to the observed SNR, SNR_{obs} , of a source, i.e. the SNR that arises from summing over the actual source mask obtained with SOFIA 2. The ratio of $\text{SNR}_{\text{obs}}/\text{SNR}_{99}$ of the sources detected in the mock data experiment in Section 6 is plotted in Fig. A1 as a function of SNR_{99} . As can be seen, the two are roughly equal at an SNR of about 30. Towards the upper end of the SNR range, SNR_{obs} becomes smaller than SNR_{99} , suggesting that SOFIA 2 is summing




Figure A1. Ratio of observed SNR and SNR_{99} as a function of SNR_{99} . The meaning of the different colours and markers is the same as in Fig. 6.

over more pixels than necessary to obtain 99 per cent of the flux of a source. The sharp rise in $\text{SNR}_{\text{obs}}/\text{SNR}_{99}$ towards the lower end of the SNR range is largely the result of noise bias and likely related to the peak flux density bias shown in the upper left panels of Figs 6 and 7.

The main conclusion from our analysis is that for faint sources the observed integrated SNR is somewhat larger than SNR_{99} , presumably due to a combination of the aforementioned noise bias and the fact that the source mask might only enclose the brighter parts of the source, thus elevating its SNR. Hence, any specific completeness limit in SNR_{99} space is likely to correspond to a somewhat higher limit in SNR_{obs} space.

In principle, it should be possible to derive a quantitative relationship between SNR_{99} and SNR_{obs} from the analysis presented here,

but in practice such a relation would be of little use in our experience, as the actual SNR_{obs} measurement will strongly depend on both the settings used in SOFIA 2 and the characteristics of the data set to be searched, and for those reasons cannot be generalized. As a basic example, choosing a higher detection threshold in the S + C finder would result in a more compact source mask and thus a larger value of SNR_{obs} , whereas enabling the mask dilation algorithm would grow the size of the mask and thus result in a smaller value of SNR_{obs} for the same source. This is because the noise scales with the square root of the number of pixels in the mask, while the signal can to first order be treated as constant.

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