



Full length article

SDHDF: A new file format for spectral-domain radio astronomy data



L.J. Toomey ^{a,*}, G. Hobbs ^a, D.C. Price ^b, J.R. Dawson ^{a,c}, T. Wenger ^{b,l}, D. Lagoy ^b,
 L. Staveley-Smith ^{d,e}, J.A. Green ^{f,j}, E. Carretti ^g, A. Hafner ^a, M. Huynh ^j, J. Kaczmarek ^{a,i},
 S. Mader ⁱ, V. McIntyre ^a, J. Reynolds ^a, T. Robishaw ^b, J. Sarkissian ⁱ, A. Thompson ^j,
 C. Tremblay ^{j,k}, A. Zic ^a

^a CSIRO Space & Astronomy, P.O. Box 76, Epping, 1710, NSW, Australia

^b International Centre for Radio Astronomy Research, Curtin University, Bentley, 6102, WA, Australia

^c School of Mathematical and Physical Sciences and MQ Research Centre in Astronomy, Astrophysics, and Astrotechnology, Macquarie University, 2109, NSW, Australia

^d International Centre for Radio Astronomy Research (ICRAR), University of Western Australia, 6009, WA, Australia

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

^f SKA Observatory, SKA-LOW Science Operations Centre, ARRC Building, 26 Dick Perry Avenue, Technology Park, Kensington, 6151, WA, Australia

^g INAF, Istituto di Radioastronomia, Via Gobetti 101, Bologna, 40129, Italy

^h Dominion Radio Astrophysical Observatory, Herzberg Astronomy and Astrophysics Research Centre, National Research Council Canada, PO Box 248, Penticton, BC V2A 6J9, Canada

ⁱ CSIRO Space & Astronomy, Parkes Observatory, Parkes, 2870, NSW, Australia

^j CSIRO Space & Astronomy, PO Box 1130, Bentley, 6102, WA, Australia

^k SETI Institute, Mountain View, CA 94043, USA

^l NSF Astronomy & Astrophysics Postdoctoral Fellow, Department of Astronomy, University of Wisconsin, Madison, WI 53706, USA

ARTICLE INFO

Keywords:

Radio astronomy

Astronomy data acquisition

ABSTRACT

Radio astronomy file formats are now required to store wide frequency bandwidths and multiple simultaneous receiver beams and must be able to account for versatile observing modes and numerous calibration strategies. The need to capture and archive high-time and high frequency-resolution data, along with the comprehensive metadata that fully describe the data, implies that a new data format and new processing software are required. This requirement is suited to a well-defined, hierarchically-structured and flexible file format. In this paper we present the Spectral-Domain Hierarchical Data Format ('SDHDF') — a new file format for radio astronomy data, in particular for single dish or beam-formed data streams. Since 2018, SDHDF has been the primary format for data products from the spectral-line and continuum observing modes at Murriyang, the CSIRO Parkes 64-metre radio telescope, and we demonstrate that this data format can also be used to store observations of pulsars and fast radio bursts.

1. Introduction

With increasing instantaneous bandwidth and higher data volumes output from modern receivers on radio telescopes, such as the Ultra-Wide-Bandwidth Low Frequency (UWL, Hobbs et al., 2020) receiver, and the cryogenically-cooled phased-array feed ('Cryo-PAF') receiver at Murriyang¹ the CSIRO Parkes 64-metre radio telescope, there was a requirement for a new file format. This requirement is not unique to

Parkes; there are many other observatories with instruments that produce multiple beams and wide bandwidths (such as the Five-hundred-meter Aperture Spherical Telescope 'FAST' in China, which utilises similar receiver technology, and has historically written data out with similar formats to Parkes). Here we present Spectral-Domain Hierarchical Data Format ('SDHDF') — a new file format for radio astronomy data based on the Hierarchical Data Format² ('HDF³').

Previously, spectral data from Murriyang were stored in 'SDFITS' format (Garwood, 2000). This format for 'Single Dish' data uses the

* Corresponding author.

E-mail address: lawrence.toomey@csiro.au (L.J. Toomey).

¹ In the Wiradjuri Dreaming, Biyami (Baiame) is a prominent creator spirit and is represented in the sky by the stars which also portray the Orion constellation. Murriyang represents the 'Skyworld' where Biyami lives.

² <https://www.hdfgroup.org/>.

³ Note: At the time of writing, any reference to HDF in this paper is to the data model of HDF version 5 (HDF5).

binary table data structure as specified in the Flexible Image Transport System ('FITS', Pence et al., 2010; Wells and Greisen, 1979). While the FITS format has enjoyed widespread usage over several decades, design decisions that were reasonable in 1980, now limit modern astronomical research; a summary of these issues is given by Thomas et al. (2014) and Thomas et al. (2015), and we discuss these later on.

The use of HDF as a file format for Parkes spectral-line and continuum data was first introduced specifically for the HI-Pulsar 'HIPSR' digital backend (Price et al., 2016) of the 21-cm multi-beam receiver, as an alternative to SDFITS. However, the HIPSR format was not intended for use outside the project, and does not implement features required for broader use cases. Amongst the wider astronomy community, HDF has been implemented at the Low Frequency Array (LOFAR, van Haarlem et al., 2013) for a decade or so (Anderson et al., 2011), albeit more recently for beam-formed data products.⁴ The Laser Interferometer Gravitational-Wave Observatory (LIGO) Scientific Collaboration publish their data in HDF (for example, the O2 Data Release⁵), and the format is also supported for image cube visualisation of large-scale data from several telescopes such as the Karl G. Jansky Very Large Array (VLA), the Atacama Large Millimeter/sub-millimeter Array (ALMA), and SKA pathfinders, with the Cube Analysis and Rendering Tool for Astronomy (CARTA.⁶) (Comrie et al., 2020)

Consequently we explored the use of HDF for all spectral-line and continuum data products from Parkes receivers, converging on the SDHDF model as the formal convention based on the following considerations:

- The storage mechanism is based on HDF, a well-defined and documented format. The HDF Group states that they are "committed to always keeping HDF software free and open source", and this is reflected in the extensive and up-to-date commit history of the repository.
- SDHDF is capable of storing comprehensive high time-resolution and wide-band high frequency-resolution data in a nested hierarchical structure suitable for telescope feeds with one beam (single pixel feed) or multiple beams, and with the potential to expand to support beam-formed data products from multiple single dishes, antennas or tiles.
- SDHDF data can be incorporated into HDF Virtual Datasets (VDS⁷), and the format is suitable for working with terabyte-scale data products in distributed computing infrastructure.
- The data and metadata are fully self-described by HDF 'attributes'.
- The format is designed to be portable and open to new functionality and capability, with long-term development support funded by donation from within the community.
- Reduction of the data products can produce secondary data products that themselves meet the SDHDF definition.
- The format can be parsed by HDF libraries available for computer languages such as C/C++, MATLAB, IDL, Julia, Python, as well as by open source tools and graphical user interfaces (GUIs).
- SDHDF data products are suitable for storage in long-term archives because the metadata are described, the files are compressible, and the SDHDF definition has been allocated a Digital Object Identifier⁸ (DOI) for the life-cycle of the data.

SDHDF was successfully commissioned as the primary format for all spectral-line and continuum data products on completion of the Parkes UWL receiver project. High time-resolution pulsar observations are still produced in 'PSRFITS' format (Hotan et al., 2004), but we demonstrate below that including pulsar data in SDHDF format would be possible in

the future. Since late 2018, more than 80 TB of SDHDF data products have been archived in the Australia Telescope Online Archive⁹ (ATOA), becoming publicly available after an initial 18 month embargo period.

In this paper, we provide an overview of the SDHDF definition in Section 2, and present design features in Section 3. In Section 4 we provide information on working with SDHDF files, including the introduction of a new software package 'INtegrated SDHDF Processing Engine in C for Telescope data Analysis' (INSPECTA), developed specifically for working with this format and for which we provide some example use cases. In Section 5 we discuss the choice of HDF as the base for our format, further SDHDF development, and uptake within the community. We present our conclusion in Section 6. An overview of the definition (Appendix A), and an example of how the format stores some observation parameters for Parkes data (Appendix B), are available in the online supplementary material¹⁰.

2. The SDHDF definition

The SDHDF definition can be thought of as a tree-like structure of HDF 'group' and 'dataset' binary objects, containing the data, observation metadata, and time, frequency and polarisation information. A visual representation of the structure is shown in Fig. 1.

SDHDF was designed for flexibility and to accommodate multi-dimensional channelised data in one or many frequency bands of any bandwidth in blocks of time of sub-second precision. A single file can contain one or many observations from different observing modes (for example a targeted pointing and/or scan) stored as multiple top-level 'root_ID' objects, each of which adhere to the SDHDF definition and can be a combination of many input data streams from one or many beams, associated metadata for both astronomy and noise source data products, and observation metadata. An SDHDF file may also be a product of post-processing from any number of other SDHDF data and metadata products structured according to the definition — these may contain more structure than the raw data. For example, we may optionally store datasets that have had mean or median averaging applied, or store additional zoomed frequency bands. The SDHDF definition.¹¹ and a template data file are publicly available on CSIRO's Data Access Portal¹²

In order to ensure that all SDHDF data products from a set of telescope observations meet the requirements of the definition, we have implemented a template methodology. With this, on completion of an observation, the data and metadata structure, descriptions, data units and values are first populated in an SDHDF template file built from the definition, where parameters are configured for a specific telescope and observing system. The HDF datasets for the astronomy data are initiated such that they are resizable. The template structure at this stage is just a placeholder — it is then populated by the raw data and metadata to form the final data product. Metadata are stored with HDF 'attributes' — objects that directly describe the dataset or group object to which they are attributed. All attributes are stored as native Python data types.

The benefits of a flexible and configurable template structure are two-fold — it facilitates uptake of the format by a range of radio astronomy institutions, and use of a template ensures that all data products from different institutions adhere to the SDHDF format definition, thereby enhancing data provenance in long-term archives.

⁴ <https://support.astron.nl/LOFARBeamformedCookbook/index.html>.

⁵ <https://doi.org/10.7935/CA75-FM95>.

⁶ <https://doi.org/10.5281/zenodo.3377984>.

⁷ <https://docs.h5py.org/en/stable/vds.html>.

⁸ <https://www.doi.org/>.

⁹ <https://atoa.atnf.csiro.au/>.

¹⁰ <https://doi.org/10.1016/j.ascom.2024.100804>

¹¹ <https://doi.org/10.25919/8807-0m95>.

¹² <https://data.csiro.au/>.

Table 1

Examples of SDHDF classes.

Class name	Description
sdhdf_file	SDHDF file class
sdhdf_band	SDHDF frequency band class
sdhdf_table	SDHDF metadata table class
sdhdf_frequency	SDHDF frequency dataset
sdhdf_data	SDHDF astronomy dataset
sdhdf_waterfall	SDHDF frequency time dataset
sdhdf_flags	SDHDF RFI flags dataset
sdhdf_weights	SDHDF RFI weights dataset

2.1. SDHDF classes and attributes

The SDHDF definition contains HDF classes and attributes that are unique to an SDHDF format file — these are useful for grouping objects together when performing operations on the data. Classes are assigned to specific dataset types, identified by the ‘SDHDF_CLASS’ attribute. Table 1 lists some examples of these classes and their descriptions.

Two classes, `sdhdf_table` and `sdhdf_waterfall` are central to the SDHDF definition. We provide a brief summary below; further details about these and other SDHDF classes can be found in the SDHDF definition. Observation metadata are stored in tables with the `SDHDF_CLASS` `sdhdf_table`. This class is similar to the HDF Table Specification Version 1.0.¹³ A one-dimensional compound HDF dataset is used to store the data, with column names ascribed in the compound datatype.

The `sdhdf_waterfall` class is an N -dimensional array with dimension scales to describe each axis. Following the HDF dimension scale specification, the attribute ‘DIMENSION_LABELS’ provides a list of strings to describe each axis. Similarly, the attribute ‘DIMENSION_LIST’ provides references to any datasets that represent a dimension scale. While the `sdhdf_waterfall` class can be used for any N -dimensional dataset, in this context the dimensions are labelled as time, product type (default: polarisation), frequency and bin,¹⁴ as defined in Section 3.1. Information about units (e.g. jansky, kelvin, counts) is stored in the ‘UNIT’ attribute.

Arrays in the `sdhdf_waterfall` datasets can be large and HDF provides the means to store multiple segments separately in the file (using a method known as ‘chunking’). The SDHDF definition does not preclude chunking any dataset in a way that will make reading and writing the data more efficient. For the Parkes observations we currently use an automatic chunking algorithm, but are currently exploring more efficient methods. We note that as long as HDF tools are used then the choice of chunking method does not affect how the data sets are loaded or processed.

Throughout the definition, other attributes are used to define special parameters — for example, the ‘FRAME’ attribute describes the spectral reference frame of the frequency dataset, set to ‘topocentric’ by default in the definition.

In Table 2 we show an example of the metadata including attributes, associated with a dataset that contains information about a telescope beam. The position of the dataset in the object hierarchy is shown in ‘Path Information’, with the data being stored as a compound entity of two different data types (‘Data Information’). Attributes of the example dataset include ‘NUMBER_OF_BANDS’ and ‘SOURCE’ entries — the definition allows for any number of beams, each with potentially different sources and associated metadata for any number of frequency bands (sub- or zoom-bands).

¹³ https://docs.hdfgroup.org/hdf5/develop/_t_b_1.html.

¹⁴ In the `sdhdf_waterfall` class, the term ‘bin’ refers to the phase of a periodic signal, such as that of a pulsar or switched noise source.

3. SDHDF design features

3.1. Flexible observing modes

The SDHDF structure is flexible enough to accommodate various observing modes. The telescope may scan across a sky region, the receiver may rotate to account for parallactic angle, or (with receivers such as phased-array feeds) individual beams may be steered independently. The top (or ‘root_ID’ level) in the SDHDF definition stores a description of each receiver beam (0 to $N-1$ beams, where ‘beam_0’ is the first beam¹⁵).

The tier down from the beam group is the ‘band’ group object. In this context a ‘band’ refers to a single range of frequencies, and may be colloquially referred to as a sub- or zoom-band. Note that we do not explicitly preclude overlapping bands; for instance a file could include a wide bandwidth and a separate band that is a sub-set of the full band. The information stored in the band group will have originally come from a signal processor, which will have formed spectra at specific time intervals.

We store time-dependent information (such as the timestamp of each spectrum, pointing information, weather conditions, etc.) independently for each telescope beam and observing band in the ‘observation_parameters’ dataset. For Parkes data we store this information with a default cadence corresponding to the rate that the time-averaged spectra are recorded, but the implementation is flexible enough to store information on any cadence. The ‘observation_parameters’ dataset is configurable based on the requirements of a particular telescope — an example of the parameters required for a Parkes observation is shown in Appendix B. Metadata information, such as the time corresponding to a specific spectrum, or the frequency of a given frequency channel, is set to the centre of the time interval or the frequency channel. The primary time information is stored as an integer Modified Julian Date (MJD) and the number of seconds (as a float) since the start of that MJD. Storing these as separate values allows the format to have sufficient time precision for pulsar data.

The spectra are stored for each beam and band in the ‘astronomy_data’ groups, as a 4-dimensional `sdhdf_waterfall` dataset with floating point precision. For the majority of continuum and spectral-line observations only a single phase bin is stored, however, for pulsar or calibration observations a large number of phase bins may be recorded. The frequency of each channel is stored in the frequency dataset. This is a two dimensional array with the dimensions being the number of integrations and the number of frequency channels. In many cases only a single integration is present (as the frequency of the channels do not change in time). However, if the frequency axis has been Doppler corrected to account for the Earth’s motion then the number of integrations likely will equal the number of spectra in the file.

The spectra recorded in the SDHDF file usually need to be calibrated. The calibration steps depend on the observing system and the particular science case. The SDHDF structure therefore needs to be sufficiently flexible to allow for different calibration methods. SDHDF was initially developed for the spectral-line and continuum output data from Murriyang, in which the astronomy spectrum is recorded with high time and frequency resolution, and spectra with the noise source switched on and separately switched off are also recorded with lower time and frequency resolution. Such information is stored in datasets for each beam and band in the ‘calibrator_data’ group with `sdhdf_waterfall` classes being used to store the on and off spectra for Parkes data. In order to confirm that the signal processor was correctly synchronised with the noise source we also used the `sdhdf_waterfall` class with multiple phase bins to record spectra as measured in 32 phase bins across the noise source period. Details of

¹⁵ In SDHDF, all indices are zero-based.

Table 2
Example metadata from an SDHDF dataset object.

Path information			
Basename	beam_parameters		
Local	sdhdf_template.hdf:/metadata/beam_parameters		
Physical	sdhdf_template.hdf:/metadata/beam_parameters		
Data information			
HDF type	COMPOUND		
Data types	LABEL (string) NUMBER_OF_BANDS (int64) SOURCE (string) RIGHT_ASCENSION (string) DECLINATION (string)		
Default attributes			
Name	Description	Unit	Value
SDHDF_CLASS	SDHDF class name	None	sdhdf_table
SDHDF_DESCRIPTION	SDHDF object description	None	Metadata specific to the antenna beam
Dimension attributes			
Name	Description	Unit	Value
LABEL	Beam label	None	beam_N
NUMBER_OF_BANDS	Number of frequency bands	None	band_N
SOURCE	Source name	None	source_name
RIGHT_ASCENSION	Coordinates in right ascension	HH:MM:SS.s	HH:MM:SS.s
DECLINATION	Coordinates in declination	DD:MM:SS.s	DD:MM:SS.s
Additional attributes			
Name	Description	Unit	Value
EQUINOX	Equinox of the coordinates	Julian year	NNNN

the calibration data can also be stored as HDF dataset attributes. For instance, the switched noise source data may be normalised so that the counts in the noise source spectra are equivalent to counts in the astronomy spectrum.

Traditional spectral-line calibration strategies rely on switching between observation parameters. For instance, the telescope position may switch between an on-source and an off-source position (or swapping beam positions), or the signal-processing system may implement a frequency-switching process. The current Parkes signal processor produces separate SDHDF data files for different source positions (i.e., on- and off-source pointings are in different files), however as shown in Fig. 1, an SDHDF file can have multiple ‘root_ID’ groups. These can be used to store multiple observations (or multiple observing setups) in a single SDHDF file. Frequency-switched observations typically shift the frequency of a local oscillator in order to facilitate spectral bandpass calibration. Within SDHDF, such observations could be stored as (1) individual SDHDF files for each change in frequency, (2) as separate ‘root_ID’ groups for each frequency, (3) by adding in an extra ‘scan’ hierarchy after the ‘band’ group, (4) by expanding the ‘observation_parameters’ metadata to include the local oscillator frequency or (5) with the inclusion of observatory-specific calibration datasets to be included in the ‘calibrator_data’ group. See Section 5.1 on extending the format.

3.2. Additional datasets, flags and weights

The SDHDF definition includes scope for extra, optional datasets. For instance, these could be required as part of the data reduction process. Perhaps the most commonly used extra datasets relate to weighting and flagging the spectra. The output of manual, or automatic, flagging routines can be stored in the ‘flags’ datasets. This is recorded with a 0 to indicate a spectral channel is not flagged and a 1 otherwise. Unless specified explicitly we use unsigned 8-bit values to record these flags to enable simple reading and writing, but they can be written as 1-bit values if required. We allow for flags to be applied to astronomy or calibration spectral channels or integrations. The PSRFITS definition does not allow flags to be applied independently for polarisation channels. However, we provide the option that flags (and weights,

as described below) can be written with the same dimensions as the raw data (including polarisation and bin information). However if, for instance, the polarisation dimension is set to 1 then the flags would be applied to all polarisations for a given time and frequency channel.

Weights can also be applied (usually to each frequency channel in each integration). The weight can represent the product of the channel bandwidth and integration time (or other weighting methods could be used such as Anderson et al., 2023). SDHDF weights are stored as floating point values in the ‘weights’ datasets.

If the data are recorded with low-bit quantisation (for example, pulsar search-mode data is generally a 2-bit data stream), then the data type is stored in the ‘DATA_TYPE’ attribute of the `sdhdf_waterfall` class. It is often necessary to provide tables of scaling factors enabling the conversion of the stored bit depth back to the measured signal strength. Following the PSRFITS definition we allow for independent `sdhdf_waterfall` datasets containing offsets and scaling parameters where the measured value v is recovered from the stored unsigned integer data value d as:

$$v = (d - z)S + \zeta \quad (1)$$

where z is a ‘zero-offset’ value, S a frequency, polarisation and time-dependent scaling factor and ζ a frequency, polarisation and time-dependent offset.

4. Working with SDHDF data

Open source GUIs and Python modules are available for interrogating HDF files in general, and the Parkes spectral-line and continuum teams have also been independently developing their own SDHDF processing tools in Python. We have developed the comprehensive INSPECTA software package, included as part of the publicly available SDHDF_TOOLS¹⁶ repository. INSPECTA is written in C and contains tools to perform a multitude of processing tasks, from visual inspection of files to calibration and even the simulation of various observing parameters. The complete list of tools is too comprehensive to list

¹⁶ https://bitbucket.csiro.au/scm/cpda/sdhdf_tools.

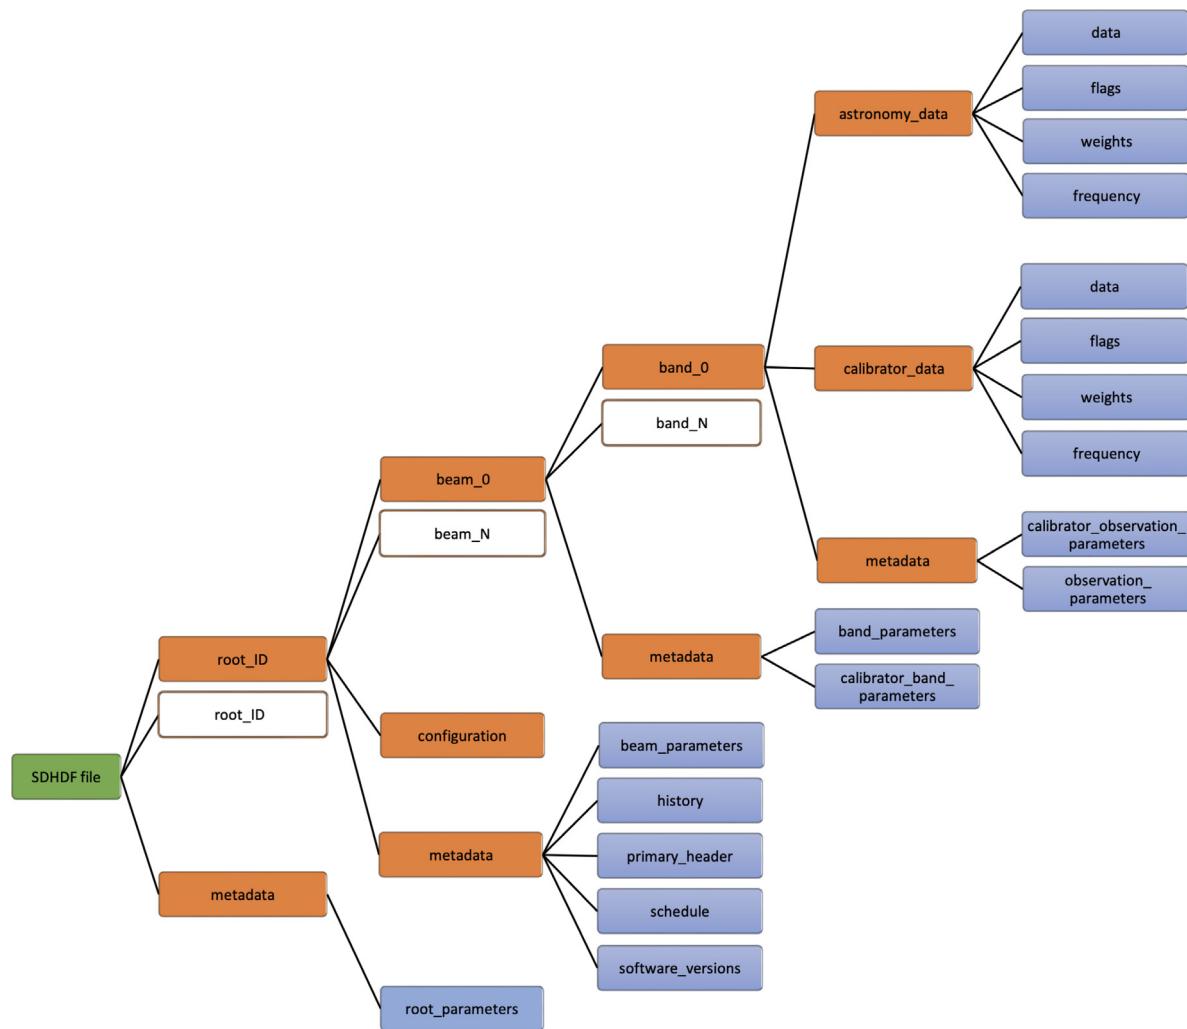


Fig. 1. A visual representation of an example data file, showing the hierarchical association of HDF groups and datasets. An SDHDF data file (the solid green rectangle) may contain optionally none or any number of ‘root_ID’ objects each of which is of the SDHDF_CLASS `sdhdf_file` and each adhering to the SDHDF definition, where each `root_ID` object could be a separate observation and grouped together (for example one may wish to store multiple single scans in the same file). Information about the `root_ID` objects are stored in a ‘root_parameters’ dataset; this dataset and the parent metadata group are not required if the file does not contain a `root_ID` object. In this Figure, the solid orange and solid blue rectangles represent HDF group and dataset objects respectively. Any number of `root_ID`, beam and band groups per beam are supported, as illustrated by the unfilled rectangles with orange borders. For Parkes UWL data, we chose to record a switching noise source for calibration purposes, storing both on, off and a 32-bin time-averaged dataset in the ‘calibrator_data’ group (not shown in this Figure).

here, but allows for spectra to be visualised, calibrated and analysed. It also allows for both automatic and manual radio-frequency interference (RFI) flagging. We provide examples below in which SDHDF files have been processed using the INSPECTA software to demonstrate the suitability of SDHDF for different types of observations. These observations have been carried out with Murriyang and the majority make use of the UWL receiver.

Here we show both the versatility of the SDHDF format and the INSPECTA software package by demonstrating how a diverse set of observations can be processed and visualised. An example is shown in the left panel of Fig. 2 in which the telescope was tracking the limb of the Moon as it occulted quasar 3C 273 on 20th February 2022 (as a modern re-enactment of Hazard et al., 1963). The SDHDF file for this observation was obtained from the ATOA.¹⁷ We recorded observations with 400 Hz resolution producing over 7 million spectral channels over the entire UWL band, which were written to disk each second. In order to ensure we do not produce a single file that is too

large to be archived, our online processing scripts split data files when they reach a specified size (currently ~10 GB) – the files can be split by beam, band or time, and the INSPECTA software can subsequently be used to join files back together as required. The total observation time for this specific data file was ~1160 s. The SDHDF file captures each of the 26 sub-bands formed by the UWL receiver system along with metadata recording the telescope pointing directions for each of the 1 s integrations. To form the output shown in Fig. 3 we used the INSPECTA tools to average the same observation in time. This is then visualised as the three RF bands for both polarisation channels. Note that no calibration has been carried out in this stage and hence the Figure shows the frequency-dependent bandpass shape and strong RFI signals primarily from mobile communication systems.

This observation can also be used to demonstrate continuum observations in which a low frequency resolution is generally required, but the data are output with short time duration, showing that the SDHDF format can store pointing information at sufficient precision and temporal resolution. We first carried out an automatic removal of both persistent and transient RFI using the INSPECTA software. We then averaged the frequency channels, summed the polarisations and plotted the total signal strength as a function of time in the 128 MHz-wide sub-band centred at 1408 MHz. Of course, we can make similar

¹⁷ <https://atoa.atnf.csiro.au/>; Project ID: PX083; File: uwl_220219_143803_3.hdf.

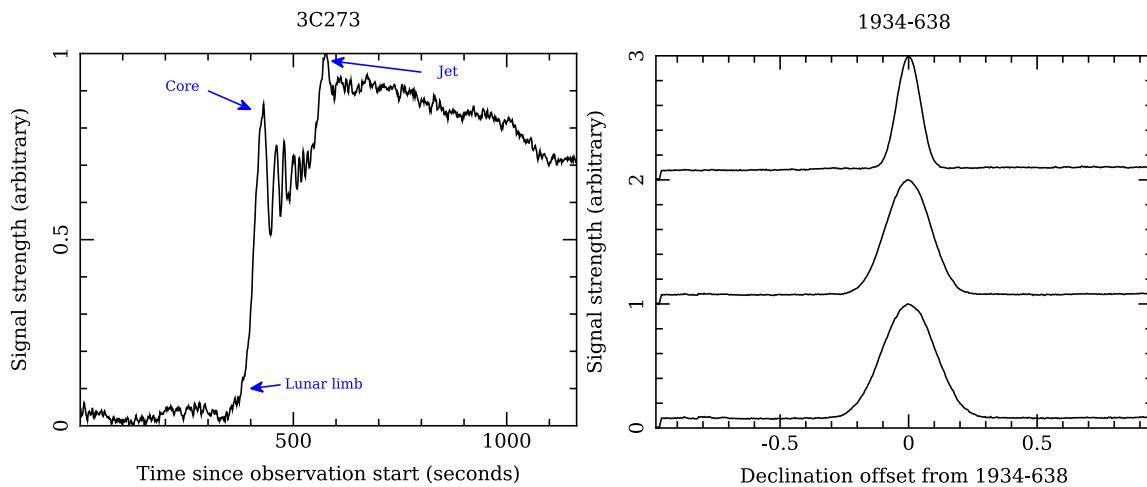


Fig. 2. (Left panel) The signal strength in the UWL 1152 MHz band as a function of time during the lunar occultation of the quasar 3C 273. The quasar core and then a jet boost the flux as they emerge from behind the lunar limb. (Right panel) Representation of a scan across a primary flux calibrator source PKS 1934–638 in three observing bands centred on 3328 (top), 1664 (centre) and 1408 MHz (bottom), where the declination offset is in degrees.

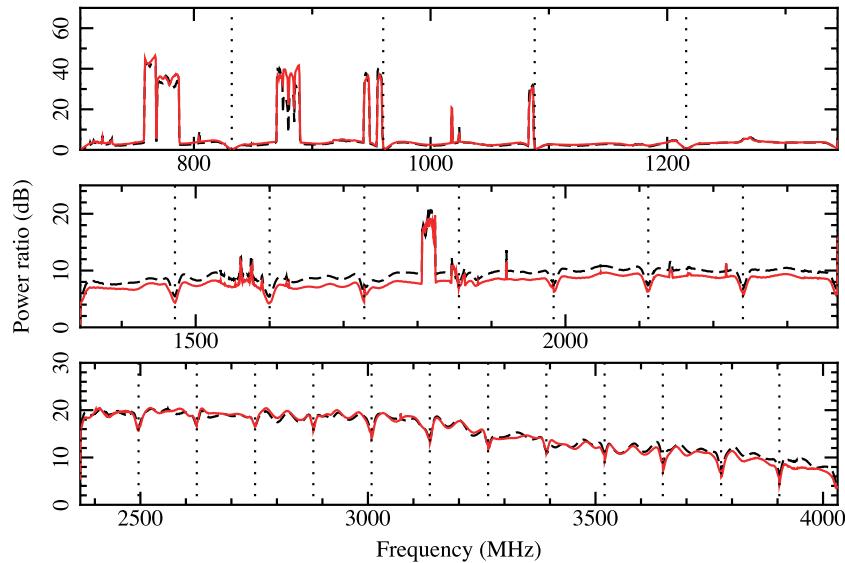


Fig. 3. A wide-bandwidth observation demonstrating the ability of SDHDF to store, process and visualise such data. The three sub-panels represent the three RF bands recorded by the Parkes UWL receiver system (704 to 1344 MHz, 1344 to 2468 MHz and 2468 to 4032 MHz), showing polarisations A and B in red and black (dashed) respectively.

plots in other sub-bands, and a full analysis of the recorded signal will be presented elsewhere.

The occultation observation was carried out by tracking the position of the source; it is more common to scan the telescope beam across a known bright source. In the right panel of Fig. 2 we show a scan across the flux density calibrator source PKS 1934–638. Here we have written three frequency bands to the SDHDF file (centred on 1408, 1664 and 3328 MHz) and the traces have been normalised and offset from each other to show the beam widths. This demonstrates that the stored metadata provide sufficient time and positional information to recover the source position. The SDHDF file for this observation was obtained from the ATOA.¹⁸

To demonstrate that SDHDF records sufficient information to calibrate the flux density and polarisation properties of a spectral-line source we use a known, bright hydroxyl maser associated with high-mass star formation, G330.953–0.182 (Caswell et al., 2014, references

therein). The SDHDF file was also obtained from the ATOA.¹⁹ For this example we do not use an off-source pointing (although the archive does contain off-source pointings obtained for this observing project). During the observation the noise source was switched on and off at a frequency of 100 Hz. The SDHDF file contains the astronomical spectrum as originally recorded, along with the spectra at lower frequency resolution when the noise source is on and also when it is off, separated out using the known phase of the noise source. We note that the integration time of the astronomy data can be different from that of the noise source. We use INSPECTA to average the data in time and to extract the part of the observing band that contains the hydroxyl lines of interest (the 1665.402 and 1667.359 MHz main-line ground-state transitions). We then use the switching noise source to calibrate the astronomy spectrum (both flux density and polarisation). This makes use of knowledge of the noise source amplitude measured through

¹⁸ <https://atoa.atnf.csiro.au/>; Project ID: P974; File: uwl_210914_053215.hdf.

¹⁹ <https://atoa.atnf.csiro.au/>; Project ID: P1073; File: uwl_220323_193627.hdf.

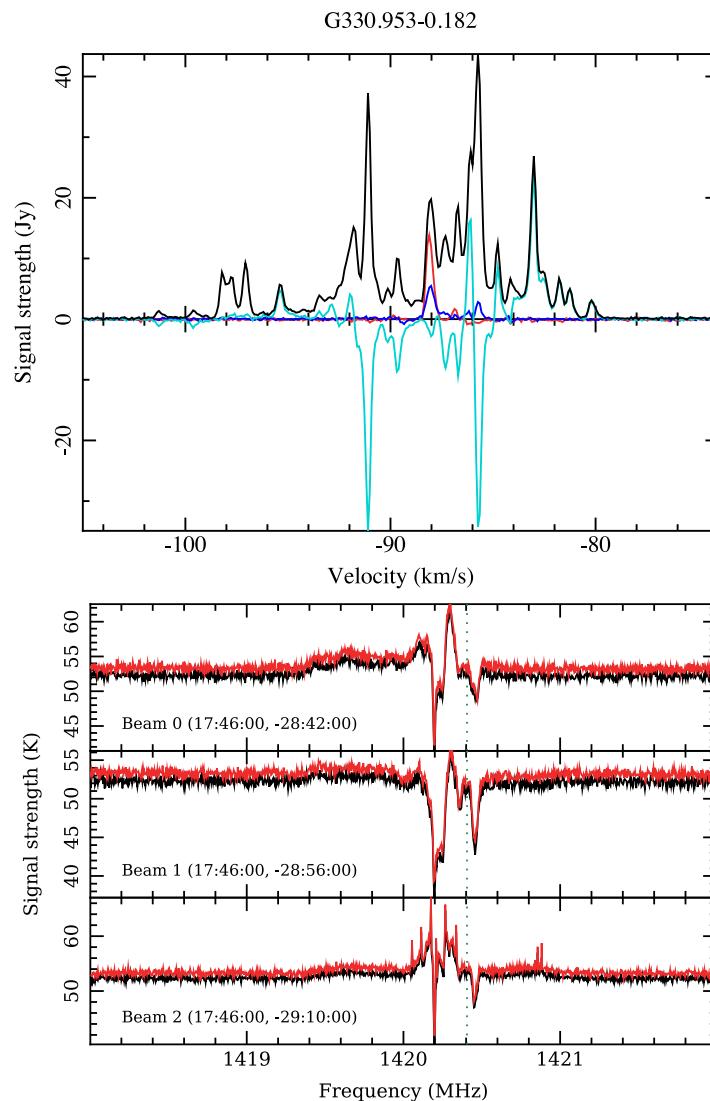


Fig. 4. (Top panel) Fully calibrated spectrum of the known 1665-MHz hydroxyl maser G330.953–0.182 from a Parkes UWL observation, showing the Stokes parameters I, Q, U and V in black, red, blue and teal respectively. (Bottom panel) Simulated example multiple-beam observations of Galactic hydrogen making use of Parkes Galactic All-Sky Survey (GASS) observations of adjacent positions separated by 14 arc-minutes (the beam full width half maximum) in declination, showing polarisations A and B in red and black respectively. Note that all beams are stored within the same SDHDF file, which can account for both fixed and moving beam directions.

separate on- and off-source observations of PKS 1934–638. These flux calibration solutions for Murriyang are available online.²⁰

The left panel in Fig. 4 shows the flux and polarisation calibrated spectrum, using the time information stored in the SDHDF file to convert the observed topocentric frequencies to the local-standard-of-rest velocity. A detailed monitoring study of the variability of this source is being carried out using similar analysis methods and will be published elsewhere.

SDHDF has been designed to allow for multiple antenna beams. However, our test case with the Parkes UWL system only provides a single beam. In order to demonstrate that archival multi-beam data can be stored in SDHDF format, and that the format is adequate for the future (such as for phased-array feeds) we have used INSPECTA to simulate data of Galactic hydrogen emission and absorption in the vicinity of the Galactic Centre based on input spectra from the Parkes Galactic All-Sky Survey (GASS; McClure-Griffiths et al., 2009) – refer to the right panel of Fig. 4. We have also demonstrated that the SDHDF

format and the INSPECTA software can be used for storing many more beams (we have successfully simulated, stored and processed 72 beams in a single SDHDF file, and the scaling is linear).

4.1. Storing phase-resolved, high time-resolution data and data precision

SDHDF was originally developed for the spectral-line and continuum radio astronomy communities. However, SDHDF can store phase-resolved data and therefore can also be used for pulsar-style observations. Additionally SDHDF has the potential to store multiple modes of observations, such as those from commensal surveys.

For nearly two decades, PSRFITS has been the standard data format for phase-resolved and time-series data such as those from pulsar fold- and search-mode observations respectively. The format is well supported by pulsar processing software packages that are trusted by the community. However, we believe that pulsar data would also benefit from being stored in SDHDF for the following reasons (noting that this is not an exhaustive list):

- PSRFITS format is not hierarchically structured, and so has no concept of the beam dimension and therefore cannot support multi-beam data.

²⁰ https://www.parkes.atnf.csiro.au/observing/Calibration_and_Data_Processing_Files.html.

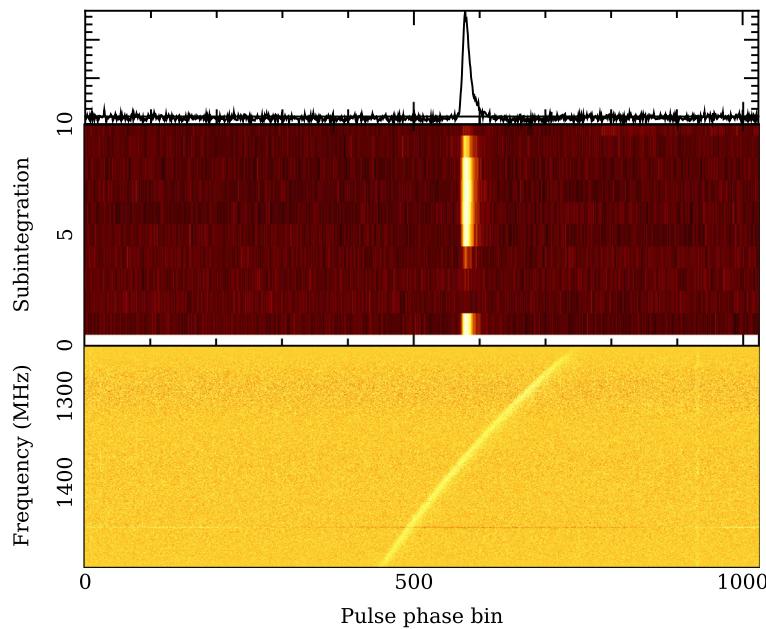


Fig. 5. Visualisation of data requiring phase bins. This observation is of the intermittent pulsar PSR J1717–4054 and the data stream has been channelised and folded online at the known period of the pulsar. The top panel shows the de-dispersed and averaged pulse profile. The central panel represents the de-dispersed pulse signal strength as a function of sub-integration, and the bottom panel shows the dispersed pulse signal strength as a function of phase and frequency.

- I/O is significantly affected if a PSRFITS binary data unit is appended to, unlike the chunked storage in HDF which allows the file to be easily extended.
- For pulsar search-mode observations, there is generally a trade-off between splitting the time-series into small chunks each with its own metadata header (as in the PSRFITS format) and having a single header and an unbroken time series (as in SIGPROC’s filterbank²¹ format). The SDHDF format allows for an unbroken time-series, whilst storing the metadata at the required cadence.
- Pulsar fold-mode observations generally start and/or end with separate calibration observations of the injected noise source, that are currently stored in separate files — these could all be incorporated into the same SDHDF file for simplicity and ease of use.
- SDHDF is able to store auxiliary data products such as images of integrated pulse profiles or gain variations which would add value to the data once archived.
- The spectral-line and pulsar software packages are independent and yet apply similar methods, for example, for calibration. A single set of algorithms can be used to ensure consistency between different observing modes.

Pulsar fold-mode data are produced by folding the incoming signal at the known period of the pulsar and storing the pulse profile for each frequency and polarisation channel with a specified number of phase bins. Such data can be stored within the SDHDF format without modification as the data structure allows for binned data. An example pulsar profile is shown in Fig. 5. The pulsar, PSR J1717–4054 was observed as part of the PULSE@Parkes outreach project (Hobbs et al., 2009) — the PSRFITS file can be obtained from the DAP.²² The file was converted into SDHDF using the INSPECTA tools and then visualised. The upper panel shows the time, polarisation and frequency averaged de-dispersed pulse profile, the central panel shows the pulse amplitude as a function of time (noting that the pulsar switches off and on during the observation) and the bottom panel shows the dispersed pulse sweep as a function of observing frequency.

High time-resolution time-series data (used, for instance, to search for pulsars or fast radio bursts) cannot be stored with floating point precision. Typically such data streams are recorded with 1, 2, 4 or 8-bits of precision. In Fig. 6 we show the first discovered fast radio burst (FRB; Lorimer et al., 2007) after converting the 13 PSRFITS data files corresponding to the 13 beams of the multi-beam receiver into a single SDHDF file. We note that the data volume in SDHDF is slightly smaller than the original PSRFITS files as we do not require the extra metadata information every 4096 samples.

5. Discussion

A commonly-held criticism and limitation of HDF is the lack of accessibility and readability — HDF is an entirely binary data format and not human-readable. Astronomers have historically enjoyed the benefits of partially text-based formats like FITS. For instance, the metadata in those files can be inspected in a standard text editor. However, we believe that the ability to manually edit a data file goes against good provenance practices, and in the Big Data era it is intractable to open an entire file in a simple text editor. We therefore believe that HDF is a suitable format for storing radio astronomy data for the following reasons:

- **Overcoming the limitations of FITS** — Existing FITS derivatives have limitations, for example RPFITS²³ did not completely adhere to the FITS standard and lacked comprehensive metadata, SD-FITS suffered from structure variation across files from different observations, and was not suitable for storing or processing of high-volume wide-frequency band data. PSRFITS, although the primary data format for pulsar astronomy data, is not flexible enough for the storage of data and metadata from spectral-line or continuum observing modes. Unlike HDF, the FITS metadata also suffer from keyword truncation, the Data Units (“extensions”) do not easily accommodate appending of large data, and there is no support for hierarchical grouping of data and metadata required by multi-beam observations.

²¹ <https://sigproc.sourceforge.net>.

²² <https://data.csiro.au/>; Project ID: P595; File: s101109_000655.rf.

²³ <https://www.atnf.csiro.au/computing/software/rpfits.html>.

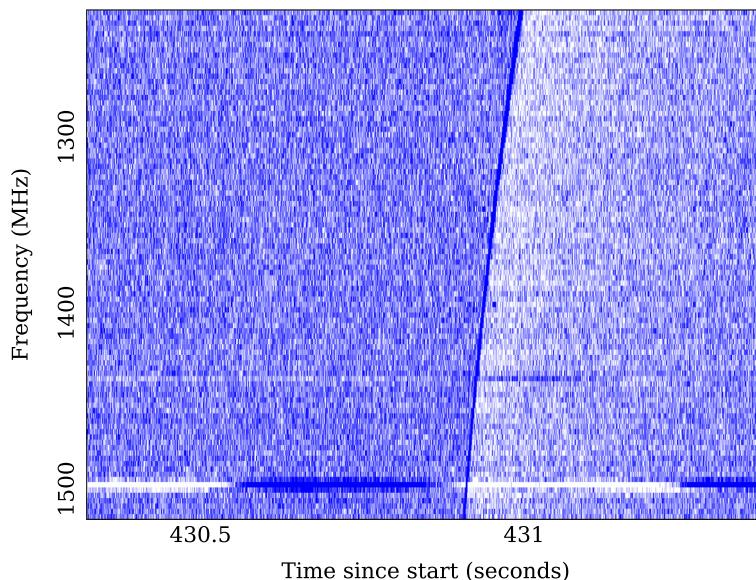


Fig. 6. High time-resolution, multi-beam data converted into SDHDF format. We show 1s of data around the first discovered fast radio burst (FRB) in beam 6 of the Parkes multi-beam receiver (Lorimer et al., 2007).

- **No suitable alternative formats** — We found no other file formats that could meet our requirements. We investigated Advanced Scientific Data Format (ASDF; Greenfield et al., 2015), a format designed to find a middle ground between FITS and hierarchically structured content, however whilst overcoming some of the limitations of FITS, there were restrictions on data and metadata storage and it was deemed too generic for storing high-volume radio astronomy data from wide-bandwidths and multiple beams.
- **Improved I/O performance and compression** — The HDF libraries support ‘chunked’ over ‘contiguous’ storage to improve I/O performance. SDHDF files can be cloud-optimised for distributed computing by writing out with increased chunk sizes. With regards to compression, historically metadata were stored sparsely, as storing multi-dimensional arrays of information such as flags would dramatically increase file size. However, the ability to compress these metadata arrays natively in HDF encourages their inclusion in the file. Price et al. (2015) demonstrated this improved I/O and compression of HDF over FITS during bench-marking of a FITS to HDF converter.
- **Accessibility** — Unlike FITS, HDF supports Unicode characters and is therefore suitable for situations where text in languages other than English is required. For instance, we have successfully demonstrated that observer names can be stored in HDF format in Chinese characters.

5.1. Extending the SDHDF format

The spectral-line and continuum astronomy communities who use Parkes for their observations are adopting the new file format and processing tools. Some teams use their own processing tools, others are now making use of INSPECTA. Developers of external tools are also adopting SDHDF. In particular AOFlagger,²⁴ an RFI flagging package, can now parse and flag SDHDF files.

The SDHDF definition is managed in an open source repository and we encourage feedback on its development. We are currently developing SDHDF version 4.0,²⁵ for which we have made improvements

to the naming scheme, dataset dimensions and attributes based on user feedback, while the underlying structure remains unchanged — important changes that apply to archival data prior to version 4.0 are shown in Tables 3 and 4 respectively.

We do expect further updates as new observing systems are commissioned and more observatories make use of this format. New datasets and attributes will be added in the future. For example, the ability to store statistics from digitisers and a flag for data quality would both be useful additions to the definition and planned for future versions. We note that SDHDF files do not require that all parameters are included. For instance, the observation metadata includes parameters such as the wind speed, which may not be relevant (or even known) for some observations. Similarly, there will always be cases in which a particular observing system requires metadata that is unique to that system. The SDHDF format does not preclude the addition of observatory-specific information.

5.2. Virtual observatory support

For the benefit of all users of astronomical data, we would like to see SDHDF become accessible and interoperable with the Virtual Observatory (VO). This can be achieved by following protocols of The International Virtual Observatory Alliance²⁶ (IVOA), an organisation that “debates and agrees the technical standards that are needed to make the VO possible”. The ATOA already has VO interoperability in place for RPFITS format data — we believe that the variety and extensive metadata included in the SDHDF format provides a significant improvement on this.

6. Conclusion

In this paper we have presented a new data format for radio astronomy data; the Spectral-Domain Hierarchical Data Format (SDHDF). The format is flexible enough to store high-time and high frequency-resolution data and metadata products from the era of large telescopes with single pixel or multi-beam receivers.

Murriyang, the CSIRO Parkes 64-metre radio telescope, is the first instrument to use the format natively for all spectral-line and continuum observations. The Ultra-Wide-Bandwidth Low Frequency receiver,

²⁴ <https://aoflagger.readthedocs.io/en/latest/>.

²⁵ We note that for data written in SDHDF < 4.0 format, users who wish to work with INSPECTA will need to checkout the ‘sdhdfv3’ branch of the software.

²⁶ <https://ivoa.net/>.

Table 3
SDHDF naming scheme changes.

SDHDF < 3.0	SDHDF 3.0 (Released: April 2023)	SDHDF 4.0 (Under development)	HDF type
config	config	Configuration	Group
cal_data_on	cal_on	calibrator_data_on	Dataset
cal_data_off	cal_off	calibrator_data_off	Dataset
cal32_data	cal_binned	calibrator_data_binned	Dataset
cal_frequency	frequency	Frequency	Dataset
data_weights	Weights	Weights	Dataset
data_flags	Flags	Flags	Dataset
obs_params	obs_params	observation_parameters	Dataset
cal_obs_params	cal_obs_params	calibrator_observation_parameters	Dataset
band_params	band_params	band_parameters	Dataset
cal_band_params	cal_band_params	calibrator_band_parameters	Dataset
beam_params	beam_params	beam_parameters	Dataset
backend_config	backend_config	instrument_configuration	Dataset
cal_backend_config	cal_backend_config	instrument_configuration	Dataset

Table 4
SDHDF dataset dimension changes.

HDF dataset	SDHDF < 3.0	SDHDF \geq 3.0	Reason
Frequency	1-dimensional	2-dimensional	Allow storage of time-dependent frequency information
Flags	2-dimensional	4-dimensional	To match the dimensions of the data, allowing flags to be applied to any dimension of the data
Weights	2-dimensional	4-dimensional	To match the dimensions of the data, allowing weights to be applied to any dimension of the data

a single pixel feed, has produced SDHDF files at Parkes since 2018, and the cryogenically-cooled phased array feed, to be commissioned in 2024, is expected to use this format for storing the data products from multiple beams.

We have developed the INSPECTA software package for the interrogation, processing and calibration of SDHDF data, and have demonstrated some of its features when working with radio astronomy data.

We envisage other institutions adopting SDHDF as their primary format because of the flexibility to store multiple data streams and comprehensive metadata, and its adherence to a formal definition suitable for long-term archiving.

CRediT authorship contribution statement

L.J. Toomey: Writing – original draft. **G. Hobbs:** Writing – original draft, Writing – review & editing. **D.C. Price:** Software. **J.R. Dawson:** Writing – review & editing. **T. Wenger:** Writing – review & editing. **D. Lagoy:** Writing – review & editing. **L. Staveley-Smith:** Writing – review & editing. **J.A. Green:** Writing – review & editing. **E. Carretti:** Writing – review & editing. **A. Hafner:** Validation. **M. Huynh:** Validation. **J. Kaczmarek:** Writing – review & editing. **S. Mader:** Writing – review & editing. **V. McIntyre:** Software, Writing – review & editing. **J. Reynolds:** Resources. **T. Robishaw:** Writing – review & editing. **J. Sarkissian:** Resources. **A. Thompson:** Software. **C. Tremblay:** Validation. **A. Zic:** Writing – review & editing.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Trey Wenger reports financial support was provided by NSF Astronomy and Astrophysics Postdoctoral Fellowship.

Data availability

All data are publicly available.

Acknowledgements

Murriyang, the CSIRO Parkes radio telescope is part of the Australia Telescope National Facility (grid.421683.a) which is funded by the Australian Government for operation as a National Facility managed by CSIRO. We acknowledge the Wiradjuri people as the Traditional Owners of the Observatory site. This paper includes archived data obtained through the Australia Telescope Online Archive (<http://atoa.atnf.csiro.au>). T.V.W. is supported by an NSF Astronomy and Astrophysics Postdoctoral Fellowship under award AST-2202340. We thank Dr. Philip Edwards for his invaluable comments on this paper.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at [doi:10.1016/j.ascom.2024.100804](https://doi.org/10.1016/j.ascom.2024.100804).

References

Anderson, K., Alexov, A., Bähren, L., Grießmeier, J.M., Wise, M., Renting, G.A., 2011. LOFAR and HDF5: Toward a new radio data standard. In: Evans, I.N., Accomazzi, A., Mink, D.J., Rots, A.H. (Eds.), *Astronomical Data Analysis Software and Systems XX*. In: *Astronomical Society of the Pacific Conference Series*, vol. 442, p. 53. [doi:10.48550/arXiv.1012.2266](https://doi.org/10.48550/arXiv.1012.2266), [arXiv:1012.2266](https://arxiv.org/abs/1012.2266).

Anderson, L.D., Liu, B., Balser, D.S., Bania, T.M., Haffner, L.M., Linville, D.J., Luisi, M., Wenger, T.V., 2023. Methods for averaging spectral line data. [doi:10.48550/arXiv.2310.09076](https://doi.org/10.48550/arXiv.2310.09076), [arXiv:2310.09076](https://arxiv.org/abs/2310.09076).

Caswell, J.L., Green, J.A., Phillips, C.J., 2014. Parkes full polarization spectra of OH masers - II. Galactic longitudes 240° to 350°. *MNRAS* 439 (2), 1680–1739. [doi:10.1093/mnras/stu046](https://doi.org/10.1093/mnras/stu046), [arXiv:1402.2726](https://arxiv.org/abs/1402.2726).

Comrie, A., Pińska, A., Simmonds, R., Taylor, A.R., 2020. Development and application of an HDF5 schema for SKA-scale image cube visualization. *Astron. Comput.* 32, 100389. [doi:10.1016/j.ascom.2020.100389](https://doi.org/10.1016/j.ascom.2020.100389).

Garwood, R.W., 2000. SDFITS: A standard for storage and interchange of single dish data. In: Manset, N., Veillet, C., Crabtree, D. (Eds.), *Astronomical Data Analysis Software and Systems IX*. In: *Astronomical Society of the Pacific Conference Series*, vol. 216, p. 243.

Greenfield, P., Droettboom, M., Bray, E., 2015. ASDF: A new data format for astronomy. *Astron. Comput.* 12, 240–251. [doi:10.1016/j.ascom.2015.06.004](https://doi.org/10.1016/j.ascom.2015.06.004).

Hazard, C., Mackey, M.B., Shimmins, A.J., 1963. Investigation of the radio source 3C 273 by the method of lunar occultations. *Nature* 197 (4872), 1037–1039. [doi:10.1038/1971037a0](https://doi.org/10.1038/1971037a0).

Hobbs, G., Hollow, R., Champion, D., Khoo, J., Yardley, D., Carr, M., Keith, M., Jenet, F., Amy, S., Burgay, M., Burke-Spolaor, S., Chapman, J., Danaia, L., Homewood, B., Kovacevic, A., Mao, M., McKinnon, D., Mulcahy, M., Osłowski, S., van Straten, W., 2009. The pulse@parkes project: a new observing technique for long-term pulsar monitoring. *PASA* 26 (4), 468–475. [doi:10.1071/AS09021](https://doi.org/10.1071/AS09021), [arXiv:0907.4847](https://arxiv.org/abs/0907.4847).

Hobbs, G., Manchester, R.N., Dunning, A., Jameson, A., Roberts, P., George, D., Green, J.A., Tuthill, J., Toomey, L., Kaczmarek, J.F., Mader, S., Marquarding, M., Ahmed, A., Amy, S.W., Bailes, M., Beresford, R., Bhat, N.D.R., Bock, D.C.J., Bourne, M., Bowen, M., Brothers, M., Cameron, A.D., Carretti, E., Carter, N., Castillo, S., Chekkala, R., Cheng, W., Chung, Y., Craig, D.A., Dai, S., Dawson, J., Dempsey, J., Doherty, P., Dong, B., Edwards, P., Ergesh, T., Gao, X., Han, J., Hayman, D., Indermuehle, B., Jeganathan, K., Johnston, S., Kanoniuk, H., Kesteven, M., Kramer, M., Leach, M., Mcintyre, V., Moss, V., Oslowski, S., Phillips, C., Pope, N., Preisig, B., Price, D., Reeves, K., Reilly, L., Reynolds, J., Robishaw, T., Roush, P., Ruckley, T., Sadler, E., Sarkissian, J., Severs, S., Shannon, R., Smart, K., Smith, M., Smith, S., Sobey, C., Staveley-Smith, L., Tzioumis, A., van Straten, W., Wang, N., Wen, L., Whiting, M., 2020. An ultra-wide bandwidth (704 to 4 032 MHz) receiver for the parkes radio telescope. *PASA* 37, e012. doi:10.1017/pasa.2020.2, arXiv: 1911.00656.

Hotan, A.W., van Straten, W., Manchester, R.N., 2004. Psrchive and psrfits: An open approach to radio pulsar data storage and analysis. *Publ. Astron. Soc. Aust.* 21 (3), 302–309. doi:10.1071/AS04022.

Lorimer, D.R., Bailes, M., McLaughlin, M.A., Narkevic, D.J., Crawford, F., 2007. A bright millisecond radio burst of extragalactic origin. *Science* 318 (5851), 777–780. doi:10.1126/science.1147532, arXiv:https://www.science.org/doi/pdf/10.1126/science.1147532, URL <https://www.science.org/doi/abs/10.1126/science.1147532>.

McClure-Griffiths, N.M., Pisano, D.J., Calabretta, M.R., Ford, H.A., Lockman, F.J., Staveley-Smith, L., Kalberla, P.M.W., Bailin, J., Dedes, L., Janowiecki, S., Gibson, B.K., Murphy, T., Nakanishi, H., Newton-McGee, K., 2009. Gass: The parkes galactic all-sky survey. I. Survey description, goals, and initial data release. *ApJS* 181 (2), 398–412. doi:10.1088/0067-0049/181/2/398, arXiv:0901.1159.

Pence, W.D., Chiappetti, L., Page, C.G., Shaw, R.A., Stobie, E., 2010. Definition of the flexible image transport system (FITS), version 3.0. *A&A* 524, A42. doi:10.1051/0004-6361/201015362.

Price, D.C., Barsdell, B.R., Greenhill, L.J., 2015. HDFITS: Porting the FITS data model to HDF5. *Astron. Comput.* 12, 212–220. doi:10.1016/j.ascom.2015.05.001, arXiv: 1505.06421.

Price, D.C., Staveley-Smith, L., Bailes, M., Carretti, E., Jameson, A., Jones, M.E., van Straten, W., Schediwy, S.W., 2016. HIPSR: A digital signal processor for the parkes 21-cm multibeam receiver. *J. Astron. Instrum.* 5 (4), 1641007. doi:10.1142/S2251171716410075, arXiv:1702.00443.

Thomas, B., Jenness, T., Economou, F., Greenfield, P., Hirst, P., Berry, D.S., Bray, E.M., Gray, N., Muna, D., Turner, J., de Val-Borro, M., Santander-Vela, J., Shupe, D., Good, J., Berrieman, G.B., 2014. Significant problems in FITS limit its use in modern astronomical research. In: Manset, N., Forshay, P. (Eds.), *Astronomical Data Analysis Software and Systems XXII*. In: *Astronomical Society of the Pacific Conference Series*, vol. 485, p. 351. doi:10.48550/arXiv.1502.05958, arXiv:1502.05958.

Thomas, B., Jenness, T., Economou, F., Greenfield, P., Hirst, P., Berry, D.S., Bray, E.M., Gray, N., Muna, D., Turner, J., de Val-Borro, M., Santander-Vela, J., Shupe, D., Good, J., Berrieman, G.B., 2014. Significant problems in FITS limit its use in modern astronomical research. In: Manset, N., Forshay, P. (Eds.), *Astronomical Data Analysis Software and Systems XXII*. In: *Astronomical Society of the Pacific Conference Series*, vol. 485, p. 351. doi:10.48550/arXiv.1502.05958, arXiv:1502.05958.

Thomas, B., Jenness, T., Economou, F., Greenfield, P., Hirst, P., Berry, D.S., Bray, E.M., Gray, N., Muna, D., Turner, J., de Val-Borro, M., Santander-Vela, J., Shupe, D., Good, J., Berrieman, G.B., Kitaeff, S., Fay, J., Laurino, O., Alexov, A., Landry, W., Masters, J., Brazier, A., Schaaf, R., Edwards, K., Redman, R.O., Marsh, T.R., Streicher, O., Norris, P., Pascual, S., Davie, M., Droettboom, M., Robitaille, T., Campana, R., Hagen, A., Hartogh, P., Klaes, D., Craig, M.W., Homeier, D., 2015. Learning from FITS: Limitations in use in modern astronomical research. *Astron. Comput.* 12, 133–145. doi:10.1016/j.ascom.2015.01.009, arXiv:1502.00996.

van Haarlem, M.P., Wise, M.W., Gunst, A.W., Heald, G., McKean, J.P., Hessels, J.W.T., de Bruyn, A.G., Nijboer, R., Swinbank, J., Fallows, R., Brentjens, M., Nelles, A., Beck, R., Falcke, H., Fender, R., Hörandel, J., Koopmans, L.V.E., Mann, G., Miley, G., Röttgering, H., Stappers, B.W., Wijers, R.A.M.J., Zaroubi, S., van den Akker, M., Alexov, A., Anderson, J., Anderson, K., van Ardenne, A., Arts, M., Asgekar, A., Avruch, I.M., Batejat, F., Bähren, L., Bell, M.E., Bell, M.R., van Bemmel, I., Bennema, P., Bentum, M.J., Bernardi, G., Best, P., Birzan, L., Bonafede, A., Boonstra, A.J., Braun, R., Bregman, J., Breitling, F., van de Brink, R.H., Broderick, J., Broekema, P.C., Brouw, W.N., Brüggen, M., Butcher, H.R., van Cappellen, W., Ciardi, B., Coenen, T., Conway, J., Coolen, A., Corstanje, A., Damstra, S., Davies, O., Deller, A.T., Dettmar, R.J., van Diepen, G., Dijkstra, K., Donker, P., Doorduin, A., Dromer, J., Drost, M., van Duin, A., Eisloffel, J., van Enst, J., Ferrari, C., Frieswijk, W., Gankema, H., Garrett, M.A., de Gasperin, F., Gerbers, M., de Geus, E., Grießmeier, J.M., Grit, T., Gruppen, P., Hamaker, J.P., Hassall, T., Hoeft, M., Holties, H.A., Horneffer, A., van der Horst, A., van Houwelingen, A., Huijgen, A., Iacobelli, M., Intema, H., Jackson, N., Jelic, V., de Jong, A., Juette, E., Kant, D., Karastergiou, A., Koers, A., Kollen, H., Kondratiev, V.I., Kooistra, E., Koopman, Y., Koster, A., Kuniyoshi, M., Kramer, M., Kuper, G., Lambropoulos, P., Law, C., van Leeuwen, J., Lemaitre, J., Loose, M., Maat, P., Macario, G., Markoff, S., Masters, J., McFadden, R.A., McKay-Bukowski, D., Meijering, H., Meulman, H., Mevius, M., Middelberg, E., Millenaar, R., Miller-Jones, J.C.A., Mohan, R.N., Mol, J.D., Morawietz, J., Morganti, R., Mulcahy, D.D., Mulder, E., Munk, H., Nieuwenhuis, L., van Nieuwpoort, R., Noordam, J.E., Norden, M., Noutsos, A., Offringa, A.R., Olofsson, H., Omar, A., Orrú, E., Overeem, R., Paas, H., Pandey-Pommier, M., Pandey, V.N., Pizzo, R., Polatidis, A., Rafferty, D., Rawlings, S., Reich, W., de Reijer, J.P., Reitsma, J., Renting, G.A., Riems, P., Rol, E., Romein, J.W., Roosjen, J., Ruiter, M., Scaife, A., van der Schaaf, K., Scheers, B., Schellart, P., Schoenmakers, A., Schoonderbeek, G., Serylak, M., Shulevski, A., Sluman, J., Smirnov, O., Sobey, C., Spreeuw, H., Steinmetz, M., Sterks, C.G.M., Stiepel, H.J., Stuurwold, K., Tagger, M., Tang, Y., Tas, C., Thomas, I., Thoudam, S., Toribio, M.C., van der Tol, B., Usov, O., van Veelen, M., van der Veen, A.J., ter Veen, S., Verbiest, J.P.W., Vermeulen, R., Vermaas, N., Vocks, C., Vogt, C., de Vos, M., van der Wal, E., van Weeren, R., Weggemans, H., Weltevrede, P., White, S., Wijnholds, S.J., Wilhelmsson, T., Wucknitz, O., Yatawatta, S., Zarka, P., Zensus, A., van Zwieten, J., 2013. LOFAR: The Low-frequency array. *A&A* 556, A2. doi:10.1051/0004-6361/201220873, arXiv:1305.3550.

Wells, D.C., Greisen, E.W., 1979. FITS - a flexible image transport system. In: Sedmak, G., Capaccioli, M., Allen, R.J. (Eds.), *Image Processing in Astronomy*. p. 445.