Astronomical Data Analysis Software and Systems XXI
ASP Conference Series, Vol. 461
Pascal Ballester, Daniel Egret, and Nuria P. F. Lorente, eds.

Status of LOFAR Data in HDF5 Format

A. Alexov,[1] P. Schellart,[2] S. ter Veen,[2] M. van den Akker,[2] L. Bähren,[1,4] J-M. Grießmeier,[3] J. W. T. Hessels,[4,1] J. D. Mol,[4] G. A. Renting,[4] J. Swinbank,[1] and M. Wise[4]

[1]*Astronomical Institute "Anton Pannekoek", University of Amsterdam, Postbus 94249, 1090 GE Amsterdam, The Netherlands*

[2]*Department of Astrophysics/IMAPP, Radboud University Nijmegen, P. O. Box 9010, 6500 GL Nijmegen, The Netherlands*

[3]*Laboratoire de Physique et Chimie de l'Environnement et de l'Espace (LPC2E), 3A Avenue de la Recherche, 45071 Orléans Cedex 2, France*

[4]*Netherlands Institute for Radio Astronomy (ASTRON), P. O. Box 2, 7990 AA Dwingeloo, The Netherlands*

Abstract. The Hierarchical Data Format, version 5 (HDF5) is a data model, library, and file format for storing and managing data. It is designed for flexible and efficient I/O and for high volume, complex data. The Low Frequency Array (LOFAR) project is solving the challenge of data size and complexity using HDF5. Most of LOFAR's standard data products will be stored using HDF5; the beam-formed time-series data and transient buffer board data have already transitioned from project-specific binary format to HDF5.

We report on our effort to pave the way towards new astronomical data encapsulation using HDF5, which can be used by future ground and space projects. The LOFAR project has formed a collaboration with NRAO, the Virtual Astronomical Observatory (VAO) and the HDF Group to obtain funding for a full-time staff member to work on documenting and developing standards for astronomical data written in HDF5. We hope our effort will enhance HDF5 visibility and usage within the community, specifically for LSST, the SKA pathfinders (ASKAP, MeerKAT, MWA, LWA), and other major new radio telescopes such as EVLA, ALMA, and eMERLIN.

1. LOFAR Introduction

The "LOw Frequency ARray" (LOFAR) is currently the largest radio telescope in the world. Almost all of the 48 stations planned are now completed and are in operation (40 in the Netherlands and 8 throughout the rest of Europe); the final Dutch stations will be available in mid-2012. A full description of LOFAR will be published soon (van Haarlem et al., in prep). LOFAR has baselines up to 1500 km and will ultimately achieve sub-arcsecond resolution over much of the the 10–240 MHz band, with a bandwidth of 48 MHz. The data is correlated on an IBM BlueGene/P supercomputer in Groningen, NL. Additional data processing and scientific pipelines are run offline, on a cluster of 100 nodes, each with 24 CPU cores, 64 GB RAM and 21 TB storage. The

Long Term Archive (LTA) has 2.2P B disk and 5 PB tape, and access to 22,600 cores via the BigGrid[1] and JUROPA[2] projects.

2. LOFAR Data Specification and Encapsulation in HDF5

LOFAR is not a traditional radio telescope: the telescope is digitally "steerable" in any direction in the visible sky; it is possible to look at multiple (up to 450) directions at the same time; any number of the 48 stations can be used individually or combined coherently or incoherently; it has additional particle detectors in the central region, called the "Superterp", which help trigger LOFAR to cosmic ray induced air shower events; and, LOFAR can observe using multiple modes in parallel, writing as many as three different data types at the same time. This complexity and flexibility creates many observing modes: examples of frequently used modes are imaging/visibility, beam-formed time-series (BF), transient buffer board dumps (TBB) and dynamic spectra. BF data and TBB data are discussed in more detail in Sections 3.1 and 3.2, respectively.

All these observing modes create data diversity and variety, and put great demands on the system: data rates up to 10 GB/sec, dataset sizes of up to hundreds of TB, and datasets with up to 6 dimensions. In order to keep up with the data rate, data writing has to be multi-threaded and parallelizable; due to I/O bottlenecks in data writing, a single LOFAR dataset is spread over multiple disks. Conventional astronomical data containers such as CASA Tables, FITS or the handful of time-series containers commonly used in Pulsar astronomy could not meet all our needs. While FITS and MeasurementSets are excellent data containers, they are limited in scope; LOFAR's data rates were a bottleneck; additionally LOFAR's new data types did not have pre-existing data models and defined formats. HDF5 is flexible and multi-purpose and can be used as the foundation for building more specific formats, like those six new formats defined by LOFAR. Using HDF5 as the foundation enables code-sharing and frees up developer time, not having to focus on low level reading, parallelization and data consistency issues, all of which have been dealt with and optimized by the HDF Group. Therefore HDF5 was deemed to be the best single file format option for encapsulating all the various LOFAR data types (see Anderson et al. (2011) and Wise et al. (2011)). The format is excellent for managing and storing large and complex scientific data.

Since 2008 the LOFAR project has been concentrating on the specification of the Interface Control Documents (ICDs) for six basic LOFAR data format types. These documents are based on scientific input on how to best store and package the data for maximum scientific clarity and usability.[3]. HDF5 allows for storage, not only of the data, but also for the associated and related meta-data describing the data's contents, conditions of observations, logs, etc. As an "all-in-one" wrapper, the HDF5 format simplifies the management of complex datasets. The LOFAR project has also developed the Data Access Library (DAL) in C++, along with an associated Python interface (PyDAL). The DAL is intermediate layer on top of HDF5; it is the software equivalent of the LOFAR ICDs for the purpose of reading, writing and accessing the LOFAR data.

---

[1] `http://www.biggrid.nl/`

[2] `http://www2.fz-juelich.de/jsc/juropa/`

[3] The LOFAR ICDs are publicly available (`http://lus.lofar.org/wiki/doku.php?id=documents:lofar_data_products`)

3. LOFAR Beam-Formed and TBB Data in HDF5

LOFAR has been taking commissioning data for over two years, for a number of modes. Most of these data were written to project-specific “bare” binary format, except for visibilities, which were written to CASA (MeasurementSet) Table format. As of 2011 LOFAR writes two of its six major data formats to HDF5: the beam-formed time-series data and the transient buffer board time-series data. These two data formats abide by the ICD specification in terms of the data structure although some of the header keyword values are not yet filled. Data are written to disk in slightly different ways, utilizing some of the unique ways that HDF5 can be used, which are discussed in more detail in Sections 3.1 and 3.2.

3.1. Beam-Formed Data in HDF5

Since the start of beam-formed time-series observations with LOFAR, data had been written in a binary format. In 2011, the LOFAR project added the capability of writing beam-formed data to both binary and HDF5 format simultaneously. The workflow is such that the LOFAR data are read from the stations by the BlueGene/P supercomputer, and processed using the online Tied-Array Beam pipeline; then the data are written to HDF5 format on the Linux cluster in two components: a binary data file and an HDF5 header “linker” file (containing the header metadata information and definition of the data structure stored within the binary file). Lastly the data are read by the science pipelines for further processing (Alexov et al. 2010). The data specifications are done with the science team within the Interface Control Document, which are also used as the specifications for the software data-writer and reader. The HDF5 linker file creates the illusion of having one HDF5 file, even though the binary file is separate. This allows us to maintain the same efficiency with data-writing without major changes to the data-writer tool suite. Having both the binary and HDF5 data on disk allows the pipeline processing software to continue running using the binary data container just as it has for two years, while at the same time allowing us to write new software to use the HDF5 file format without interrupting the data processing flow. We have been successful at writing up to 7.8 TB of data for a single observation, where files are spread over several cluster nodes, using the binary plus HDF5 linker file method.

3.2. TBB Data in HDF5

Each of LOFAR’s 48 stations consists of hundreds of dipole antennas; each LOFAR dipole antenna has a ring buffer, transient buffer board (TBB), that can store up to 1.3 seconds of raw time-series data sampled every 5 ns (LOFAR’s highest time resolution). This mechanism allows the data to be buffered after digitization. With these TBB boards LOFAR acts as a “time machine”—a transient event triggers the buffered data to be written to HDF5, providing historical information with 100 MHz bandwidth and full temporal and spatial resolution: this is a “TBB dump”. These dumps are stored in HDF5, via the DAL software layer, abiding by the ICD specification.[4]

Note that the setup is flexible and that the look-back time can be 13 seconds at 10 MHz bandwidth. The TBB datasets vary in size from 40 MB to 48 GB per station,

---

[4]The TBB ICD is available on the LOFAR data products documentation page: `http://lus.lofar.org/wiki/doku.php?id=documents:lofar_data_products`

which equates to 2 GB–2.25 TB total dataset sizes, depending on the number of stations used and the requested duration-time. In 2012, the TBB memory is likely to be upgraded, allowing 5–50 seconds of look-back time, for 100 MHz or 10 MHz respectively, pushing the data size to a maximum of 8.6 TB for one TBB dump. Phenomena that can be studied using this LOFAR functionality include air showers produced by cosmic rays, lightning on Earth and other planetary bodies within our own solar system and unknown sub-second timescale astronomical transients.

4. Summary, Collaborations And Future Considerations

We believe that the HDF5 format is well-suited for complex and large astronomical data. For some of LOFAR's data format challenges HDF5 is the only solution which provides the performance and flexibility required to fully exploit the unique capabilities of the LOFAR telescope. In HDF5, LOFAR can easily map tens of TB of complex, 6-dimensional data structures with intricate header and logging information.

The LOFAR development team has started work on the next generation of its C++ Data Access Library (DAL + PyDAL), based on current time-series data use cases and use patterns, with the goal of speed-up, simplification of use and better python bindings. We also intend to expand the DAL for use with additional LOFAR data formats and to write converters to/from FITS, CASA Tables and HDF5.

LOFAR has teamed up with NRAO, the VAO and The HDF Group in requesting funding for defining and developing astronomical data standards in HDF5, based on LOFAR's ICD and DAL ground work. This effort is nicknamed "AstroHDF". More information can be found in Masters et al. (2012) and Wise et al. (2011).

We feel that HDF5 can meet the data needs and demands of future projects such as LSST, JWST and the SKA, where size and complexity will be an even greater challenge than that faced by LOFAR. To engage in the ongoing discussion and collaboration on using HDF5 for astronomical data encapsulation, the nextgen-astrodata@astron.nl mailing list has been changed to the Google Group "nextgen-astrodata". **Please sign up on Google (`http://groups.google.com/group/nextgen-astrodata/`)** to keep the conversation flowing on how to solve the issues of large and complex data in astronomy.

References

Alexov, A., Hessels, J., Mol, J. D., Stappers, B., & van Leeuwen, J. 2010, in ADASS XIX, edited by Y. Mizumoto, K.-I. Morita, & M. Ohishi, vol. 434 of ASP Conf. Ser., 193. 1012.1583

Anderson, K., Alexov, A., Bähren, L., Grießmeier, J.-M., Wise, M., & Renting, G. A. 2011, in ADASS XX, edited by I. N. Evans, A. Accomazzi, D. J. Mink & A. H. Rots, vol. 442 of ASP Conf. Ser., 53. 1012.2266

Masters, J., Alexov, A., Folk, M., Hanisch, R., Heber, G., & Wise, M. 2012, in ADASS XXI, edited by P. Ballester, D. Egret, & N. P. F. Lorente, vol. 461 of ASP Conf. Ser., 871

Wise, M., Alexov, A., Folk, M., Pierfederici, F., Anderson, K., & Bähren, L. 2011, in ADASS XX, edited by I. N. Evans, A. Accomazzi, D. J. Mink & A. H. Rots, vol. 442 of ASP Conf. Ser., 663