



Publication Year	2021
Acceptance in OA	2022-03-22T11:01:48Z
Title	A Real Time Processing system for big data in astronomy: Applications to HERA
Authors	La Plante, P., Williams, P. K. G., Kolopanis, M., Dillon, J. S., Beardsley, A. P., Kern, N. S., Wilensky, M., Ali, Z. S., Abdurashidova, Z., Aguirre, J. E., Alexander, P., Balfour, Y., BERNARDI, GIANNI, Billings, T. S., Bowman, J. D., Bradley, R. F., Bull, P., Burba, J., Carey, S., Carilli, C. L., Cheng, C., DeBoer, D. R., Dexter, M., de Lera Acedo, E., Ely, J., Ewall-Wice, A., Fagnoni, N., Fritz R., Furlanetto, S. R., Gale-Sides, K., Glendenning, B., Gorthi, D., Greig, B., Grobbelaar, J., Halday, Z., Hazelton, B. J., Hewitt, J. N., Hickish, J., Jacobs, D. C., Julius, A., Kerrigan, J., Kittiwisit, P., Kohn, S. A., Lanman, A., Lecalake, T., Lewis, D., Liu, A., MacMahon, D., Malan, L., Malgas, C., Maree, M., Martinot, Z. E., Matsetela, E., Mesinger, A., Molewa, M., Morales, M. F., Mosiane, T., Murray, S., Neben, A. R., Nikolic, B., Parsons, A. R., Pascua, R., Patra, N., Pieterse, S., Pober, J. C., Razavi-Ghods, N., Ringuette, J., Robnett, J., Rosie, K., Santos, M. G., Sims, P., Smith, C., Syce, A., Thyagarajan, N., Zheng, H.
Publisher's version (DOI)	10.1016/j.ascom.2021.100489
Handle	http://hdl.handle.net/20.500.12386/31770
Journal	ASTRONOMY AND COMPUTING
Volume	36

A Real Time Processing System for Big Data in Astronomy: Applications to HERA

Paul La Plante^{a,b,c,*}, Peter K. G. Williams^{d,e}, Matthew Kolopanis^f, Joshua S. Dillon^{a,1}, Adam P. Beardsley^{f,g,1}, Nicholas S. Kern^{a,h}, Michael Wilenskyⁱ, Zaki S. Ali^a, Zara Abdurashidova^a, James E. Aguirre^c, Paul Alexander^j, Yanga Balfour^k, Gianni Bernardi^{l,m,k}, Tashalee S. Billings^c, Judd D. Bowman^f, Richard F. Bradleyⁿ, Phil Bull^o, Jacob Burba^p, Steve Carey^j, Chris L. Carilli^q, Carina Cheng^a, David R. DeBoer^a, Matt Dexter^a, Eloy de Lera Acedo^j, John Ely^j, Aaron Ewall-Wice^a, Nicolas Fagnoni^j, Randall Fritz^k, Steven R. Furlanetto^r, Kingsley Gale-Sides^j, Brian Glendenning^q, Deepthi Gorthi^a, Bradley Greig^s, Jasper Grobbelaar^k, Ziyaad Halday^k, Bryna J. Hazelton^{i,t}, Jacqueline N. Hewitt^h, Jack Hickish^a, Daniel C. Jacobs^f, Austin Julius^k, Joshua Kerrigan^p, Piyanat Kittiwisit^u, Saul A. Kohn^c, Adam Lanman^p, Telalo Lekalake^k, David Lewis^f, Adrian Liu^v, David MacMahon^a, Lourence Malan^k, Cresshim Malgas^k, Matthys Maree^k, Zachary E. Martinot^c, Eunice Matsetela^k, Andrei Mesinger^w, Mathakane Molewa^k, Miguel F. Moralesⁱ, Tshagofalang Mosiane^k, Steven Murray^f, Abraham R. Neben^h, Bojan Nikolic^j, Aaron R. Parsons^a, Robert Pascua^{a,v}, Nipanjana Patra^a, Samantha Pieterse^k, Jonathan C. Pober^p, Nima Razavi-Ghods^j, Jon Ringuetteⁱ, James Robnett^q, Kathryn Rosie^k, Mario G. Santos^{k,u}, Peter Sims^p, Craig Smith^k, Angelo Syce^k, Nithyanandan Thyagarajan^{q,2}, Haoxuan Zheng^h

^aDepartment of Astronomy, University of California, Berkeley, CA

^bBerkeley Center for Cosmological Physics, University of California, Berkeley, CA

^cDepartment of Physics and Astronomy, University of Pennsylvania, Philadelphia, PA

^dCenter for Astrophysics | Harvard & Smithsonian, Cambridge, MA

^eAmerican Astronomical Society, Washington, DC

^fSchool of Earth and Space Exploration, Arizona State University, Tempe, AZ

^gPhysics Department, Winona State University, Winona, MN

^hDepartment of Physics, Massachusetts Institute of Technology, Cambridge, MA

ⁱDepartment of Physics, University of Washington, Seattle, WA

^jCavendish Astrophysics, University of Cambridge, Cambridge, UK

^kSouth African Radio Astronomy Observatory, Cape Town, South Africa

^lDepartment of Physics and Electronics, Rhodes University, Grahamstown, South Africa

^mINAF-Istituto di Radioastronomia, Bologna, Italy

ⁿNational Radio Astronomy Observatory, Charlottesville, VA

^oQueen Mary University London, London, UK

^pDepartment of Physics, Brown University, Providence, RI

^qNational Radio Astronomy Observatory, Socorro, NM

^rDepartment of Physics and Astronomy, University of California, Los Angeles, CA

^sDepartment of Physics, University of Melbourne, Parkville, VIC 2010, Australia

^teScience Institute, University of Washington, Seattle, WA

^uDepartment of Physics and Astronomy, University of Western Cape, Cape Town, South Africa

^vDepartment of Physics and McGill Space Institute, McGill University, Montreal, Canada

^wScuola Normale Superiore, 56126 Pisa, PI, Italy

Abstract

As current- and next-generation astronomical instruments come online, they will generate an unprecedented deluge of data. Analyzing these data in real time presents unique conceptual and computational challenges, and their long-term storage and archiving is scientifically essential for generating reliable, reproducible results. We present here the real-time processing (RTP) system for the Hydrogen Epoch of Reionization Array (HERA), a radio interferometer endeavoring to provide the first detection of the highly redshifted 21 cm signal from Cosmic Dawn and the Epoch of Reionization by an interferometer. The RTP system consists of analysis routines run on raw data shortly after they are acquired, such as calibration and detection of radio-frequency interference (RFI) events. RTP works closely with the Librarian, the HERA data storage and transfer manager which automatically ingests data and transfers copies to other clusters for post-processing analysis. Both the RTP system and the Librarian are public and open source software, which allows for them to be modified for use in other scientific collaborations. When fully constructed, HERA is projected to generate over 50 terabytes (TB) of data each night, and the RTP system enables the successful scientific analysis of these data.

Keywords: methods: data analysis — physical sciences and engineering: astronomy — software: data analysis — software: development

1. Introduction

In recent years, the amount of scientific data has exploded. Astronomy is no exception to this transformation, and experiments are generating more data than at any point in the past. The acceleration of data generation has outpaced Moore’s Law for storage, causing data *transfer* to join data *storage* among the technical challenges that researchers must grapple with. These closely related computational requirements of analyzing data and storing them are near-universal problems that must be overcome to enable the scientific goals of these experiments. Indeed, as the size and number of files continues to grow, the very process of handling the data becomes non-trivial to solve and can rival the sophistication of the actual analysis being performed.

The field of radio astronomy boasts some of the highest data rates and volumes in all of astronomy and astrophysics research. Digitization of the radio spectrum (as a function of time, frequency, and instrumental polarization component) is close to an “embarrassingly parallel” problem, so that the data rate out of a radio telescope essentially is only bounded by its budget for digital signal processing (DSP) hardware. Furthermore, many radio telescopes are interferometers, which cross-correlate the signals measured by pairs of receivers, yielding a total data rate that scales as the square of the number of receivers. The computer clusters that perform interferometric cross-correlations are among the largest single-purpose machines built for scientific research and have some of the highest sustained-throughput data rates in the world. Radio observatories aim to operate with high duty cycles, generally 12–24 hours a day, such that their large data rates lead to large data volumes as well.

Once data have been taken, compute-intensive tasks such as calibration and the excision of radio frequency interference (RFI) generally must be performed before scientific analysis can proceed. Due to the high duty cycles of typical radio observatories, these algorithms must run reliably in close to real time, a substantial challenge at the data rates of current-generation experiments such as the Hydrogen Epoch of Reionization Array (HERA³, DeBoer et al. (2017)), let alone next-generation ones such as the Square Kilometre Array (SKA⁴). The need for real-time analysis is further enhanced because the data rates of these experiments are such that it is computationally or financially unfeasible to store the full raw data. Immediate data reduction is necessary to attain data volumes that can be transferred or archived.

In order to deal with the ever-increasing demands of data storage and processing, the HERA collaboration has

developed the Real Time Processing (RTP) system and the Librarian, which together support the analysis and storage of raw and reduced HERA data products. The RTP system is built on a Python package called `hera_opm`⁵ (ascl:2104.001)—the HERA Online Processing Module—which defines and manages a workflow of analysis steps. At the same time, RTP is also composed of various monitoring systems that track the progress of analysis steps, with an eye toward automatic diagnosis of potential processing issues. The Librarian⁶ (ascl:2104.002) is also a Python package and supports long-term storage of raw and processed data files and facilitates movement of data between different computing facilities. Both of these packages are publicly available and licensed under open-source licenses, with the hope that they may be widely useful to the broader astronomical research community and beyond.

Although in this paper we discuss these systems primarily in the context of HERA, the frameworks are sufficiently general that they may be adapted to other purposes with relatively little modification. For example, the Librarian has been adopted by the Simons Observatory (SO⁷), and the primary functionality of `hera_opm` does not reference specifics of HERA data files or analysis techniques. The RTP and the Librarian systems have been built to be modular, which allows for them to be adopted to use in other contexts. At the same time, their well-documented and reliably tested codebase can provide a stable platform on which to build, without the need to reinvent existing infrastructure. The paper below is organized as follows: in Sec. 2, we outline the data processing and storage requirements relevant for HERA. In Sec. 3, we describe the RTP system and present examples of its use. In Sec. 4, we describe the Librarian system. In Sec. 5 we provide additional discussion and describe future directions of these systems.

2. Processing and Storage Requirements

In principle, an interferometer can measure a full correlation matrix of cross-correlated antenna signals for a series of times and frequencies. This is a significant amount of data which must be moved, calibrated, imaged, etc. Although in practice there are often cuts made to the number of antennas, times, and/or frequencies used in data analysis schemes, maximizing the sensitivity and scientific return of an experiment necessitates handling as much of the data as possible in a reliable and efficient manner.

An array with N_{ant} elements has $N_{\text{ant}}(N_{\text{ant}} + 1)/2$ unique pairs of antennas (including auto-correlations). Each of these baselines produces a spectrum consisting of N_{freq} frequency channels, which is produced by the correlator every Δt_{corr} seconds. Such a spectrum is produced for each polarization element of the antenna, and

*Corresponding Author

Email address: plaplant@berkeley.edu (Paul La Plante)

¹NSF Astronomy and Astrophysics Postdoctoral Fellow

²NRAO Jansky Fellow

³<https://reionization.org>

⁴<https://skatelescope.org>

⁵https://github.com/HERA-Team/hera_opm

⁶<https://github.com/HERA-Team/librarian>

⁷<https://simonsobservatory.org>

cross-correlated when forming visibilities to produce N_{pol} instrumental polarization spectra. For a given observation window T_{obs} , the total number of spectra generated is $N_{\text{time}} = T_{\text{obs}}/\Delta t_{\text{corr}}$. Visibilities are typically recorded as single-precision floating-point complex numbers, where the real and imaginary components each require 4 bytes to store. Thus, the total volume of data V produced by HERA in a single night of observation is:

$$V = \frac{1}{2} N_{\text{ant}} (N_{\text{ant}} + 1) N_{\text{freq}} N_{\text{pol}} N_{\text{time}} \times 2 \times 4 \text{ bytes.} \quad (1)$$

The fully constructed HERA array will consist of $N_{\text{ant}} = 350$, $N_{\text{freq}} = 6144$, $N_{\text{pol}} = 4$, and produce spectra every 2 seconds. For a typical observation of 12 hours ($N_{\text{time}} = 21,600$), this leads to a raw data volume of over 260 terabytes (TB). The actual data volume recorded by HERA is significantly smaller than this (e.g., using baseline dependent averaging, as described in Wijnholds et al. 2018), though the resulting data volume per night is projected to be well over 50 TB. A typical observing season for HERA is 100 days, so several petabytes of data are generated within a single observing season.

Like many other radio observatories, HERA is situated in a remote location to avoid terrestrial sources of RFI. In this case, HERA is in the Karoo desert of South Africa, hosted by the South African Radio Astronomy Observatory (SARAO⁸). This means that both on-site computational resources and network bandwidth to the wider world are limited, raising further challenges. A representative bandwidth of 100 Mbps is only capable of moving about 1 TB per day, and so significant processing of raw data must be done to facilitate adequately small data products. Because HERA is designed to observe at a $\sim 50\%$ duty cycle (12 hours every night), the RTP system must be able to do its work within 24 hours to avoid falling behind on data analysis. This requirement adds time pressure to the RTP system’s processing requirement, as the raw data cannot be stored indefinitely. Accordingly the RTP system monitors files being processed and identifies potential blockages of the pipeline. Finally, data products of sufficient quality must be produced to allow HERA to achieve its science goals, which requires a dynamic range between foreground signals and the expected cosmological signal of $\sim 10^5$ (Pober et al., 2013).

2.1. Design Considerations

Before discussing the RTP and the Librarian systems in detail, we begin by outlining some of the basic requirements of the systems, and some of the considerations driving the design of these systems. By laying out these ideas, we motivate the choices made in the systems, and demonstrate the applicability to other systems with similar applications.

For the RTP system, the main features are:

- enabling a way to easily define a workflow (i.e., fixed set of tasks) and apply it to an arbitrary set of input data;
- allowing for flexibly interfacing with various cluster resource schedulers (e.g., Slurm,⁹ TORQUE,¹⁰ etc.);
- providing the option for the user to group together and operate on “sets” of files if desired (e.g., flagging of radio-frequency interference (RFI) may work better on longer “chunks” of data);
- ensuring high uptime and communication with users about the status of the pipeline execution.

For the Librarian system, the main features are:

- providing a reliable way to save telescope data for long-term storage;
- communicating between multiple Librarian servers and transferring data easily;
- ensuring data integrity when replicating data across servers;
- allowing for adding and managing multiple storage locations of a Librarian server;
- supporting access by automated pipelines;
- automating data movement within and between sites.

For both systems, we also want to be running based on a highly tested and well-documented code base to provide confidence in scientific results. This approach helps build confidence that the results are reliable.

Although the RTP and the Librarian have been developed expressly in the context of HERA, there are other data management and transport systems that are used by other observatories and experiments. One such system is the Next Generation Archive System (NGAS¹¹), which is used for other radio astronomy experiments such as the Atacama Large Millimeter Array (ALMA¹²). The Murchison Widefield Array (MWA¹³) uses a system known as Manta-ray¹⁴, which serves as a client for the data archiving and storage system. The Low Frequency Array (LOFAR¹⁵) hosts a “Long Term Archive” of 43 petabytes (PB) which supports archiving and retrieval of data for scientific analysis (Spreeuw et al., 2019). Outside of the field of radio astronomy, large collaborations like the Large Hadron Collider have developed systems for handling enormous data

⁸<https://www.sarao.ac.za/>

⁹<https://slurm.schedmd.com/>

¹⁰<https://adaptivecomputing.com/cherry-services/torque-resource-manager/>

¹¹<https://github.com/ICRAR/ngas>

¹²<https://www.almaobservatory.org/en/home/>

¹³<https://www.mwatelescope.org>

¹⁴<https://github.com/MWATelescope/manta-ray-client>

¹⁵<https://www.astron.nl/telescopes/lofar>

volumes, and devised plans for dealing with these issues in future research endeavors (Albrecht et al., 2017). However, none of these solutions quite fit the use needs of the HERA collaboration: the source code was sometimes proprietary, and installation was often difficult. Other solutions are primarily for a long-term archive system that does not necessarily support frequent access for real-time analysis. Thus, the need arose to build a series of tightly coupled data-processing and data-storage systems to facilitate the needs of the collaboration.

In addition, HERA is an official pathfinder of the SKA, whose data processing and handling needs will be even more demanding than those of HERA. Although there is no formal agreement between the collaborations, many of the developments made throughout the course of the experiment can be used to inform design decisions for the SKA. Additionally, having the RTP and Librarian code bases as public repositories with documentation can provide important insights to be used as more of the SKA system is built and developed.

2.2. Terminology

Throughout the rest of the discussion, we make use of several key words that have special meaning to the RTP and the Librarian systems. We briefly summarize these terms here, so that the reader may refer back to them later.

- *Observation*: a single file generated by the HERA correlator system. Due to internal throughput constraints of the correlator, these tend to be 16 seconds in length, and are projected to be ~ 25 GB in size when observing for 350 antennas.
- *Observing session*: the combined data product for a continuous set of *observations*. Generally for HERA these comprise 10-12 hours of sustained observation. They are composed of roughly 2000 individual observations.
- *File*: in the context of the Librarian, this is an abstract definition of any data product. While it may correspond to a single file in the traditional “filesystem-like” sense, it may also be a directory treated as a single object. For data generated by telescopes, they can be indexed by membership in an *observation* or an *observing session*. As discussed more below in Sec. 4.1, the name, size, and hash are required/guaranteed to be unique for a particular data product.
- *File instance*: in the context of the Librarian, this is a specific *file* that is stored on a particular Librarian server. As discussed more below in Sec. 4.1, a Librarian server is allowed to have multiple copies of a given file instance, or may not have any local instances at all.

In the following discussion, we have emphasized these words when used in a context that connotes the above meanings.

3. Real Time Processing

The real time pipeline developed for the reduction of HERA data has emerged after several iterations stretching back over years of development beginning during the PAPER project.¹⁶ In these various iterations we have investigated ways to support a range of operations with notable data processing needs. What emerged is a global system for managing large data volumes (Librarian), a processing system for running workflows on data (RTP), and a monitor and control system. Here we focus primarily on RTP but provide a brief description of the other supporting elements.

Though the discussion here is specific to HERA, many of the tools developed and employed are publicly available and applicable to current and future observatories with intense data processing needs. By presenting the general functionality of these systems, we hope that other collaborations may find these tools useful for constructing real-time systems for data-intensive tasks.

The focus of this section is primarily on the software infrastructure developed for running HERA data analysis. We do not focus as much on the actual analysis steps being run, which include RFI excision and calibration of visibilities. For a more in-depth discussion of these analysis steps, see Dillon et al. (2020); Kern et al. (2020a,b); The HERA Collaboration et al. (2021); Aguirre et al. (2021). Though there is some discussion of the specific processing steps, the architecture presented is sufficiently flexible to be adapted to other applications in a relatively straightforward fashion.

There are several existing pipeline management packages that already exist (such as SciLuigi,¹⁷ Jug,¹⁸ and COSMOS¹⁹) so it is worth discussing what features the RTP system provides that existing infrastructure does not. The primary use-case this package addresses that others do not is the ability to apply a series of steps to a list of input data. Specifically, this applies to data files that are time-ordered, with “neighbors” that come either before or after them. These individual files can have arbitrary time boundaries, which can change from night-to-night depending on the start- and stop-times of the observation. Although efforts are made to be aligned to a common grid in local sidereal time (LST) to facilitate averaging data together from different nights, the precise start and stop times otherwise can be arbitrary. For instance, these times may change based on the time of year to avoid observing

¹⁶One such example is the RTP repository, though the original codebase is no longer under active development.

¹⁷<https://github.com/pharmbio/sciluigi>

¹⁸<https://github.com/luispedro/jug>

¹⁹<https://mizzou-cbmi.github.io/>

the sun, which can lead to systematic errors in data analysis and temperature-related fluctuations in the behavior of electronics. Many algorithms, such as a time domain convolution, are applied on a time range larger than the native length of a data file. RTP handles the concept of pre-requisite jobs in a user-friendly manner. These pre-requisites specify tasks that must complete before a given task is launched, usually because its output is required for a future step. These pre-requisites can include a prior step in the workflow for the file itself, or for an arbitrary number of time-neighbors either before or after it. In addition to allowing the user to specify pre-requisites in time, the user is also able to provide a partitioning of input files along the time axis to pass to a specific task. This allows the user flexibility in handling file processing for different steps with different requirements, where some steps may require a particular number of input files from multiple times, but other steps should operate on each file individually.

We discuss these features in more detail below. In Sec. 3.1, we discuss the general RTP architecture and designed use-case. In Sec. 3.2, we talk about its use in the real-time system employed by HERA. In Sec. 3.3, we discuss additional workflows implemented using the `hera_opm` framework in service of HERA data processing, and outline some of the flexibility offered by the package. In Sec. 3.4, we briefly discuss some of the monitoring tools that the RTP system uses to keep track of the status of processing.

3.1. RTP Architecture

The RTP system oversees the successful execution of the HERA analysis pipeline. RTP consists of a loosely coupled set of independent processes, as well as monitoring daemons that track the progress of analysis in the system as a whole. These monitoring systems interface with other systems on site (such as the Monitoring & Control system²⁰). Many of the tools that it relies on feature open-source licenses and are freely available to download. Most of the core functionality is written in Python, which makes the system highly portable. Additionally, many of these tools are industry-standard. Typical clusters may already have them available to users or they can be installed with little difficulty.

Once raw data are recorded by the correlator, the RTP system begins operating. The raw data storage nodes are attached to the processing nodes via network file system (NFS) mounted directories, which allows for accessing data without explicitly copying files to local scratch space. Output files are also written to this shared disk space. The analysis pipeline launches automatically at the conclusion of observing each night without the need of human intervention. Automation helps ensure that all data can be examined within the time constraint, and so the RTP system automatically retries jobs in an attempt to provide some robustness to occasional failures.

The primary goal of RTP is to run a series of compute-intensive analysis tasks. For the workflow execution, it uses the `hera_opm` package. This package converts a user-defined workflow file (specified in the TOML format²¹, an example of which is shown in Listing 1) and a series of input files into concrete steps in an overall work diagram. This diagram is written in a file format that can be parsed by the `makeflow` package, which is available as part of the `cctools`²² system (Carmichael et al., 2010; Albrecht et al., 2012). Internally, `makeflow` generates a directed acyclic graph given dependencies for each step in the process, similar to the `make` command line utility. Once this graph has been generated, it manages the execution of each step in the process either locally or using cluster scheduling systems. When using cluster management systems, `makeflow` supports specifying batch job options, such as the number of processing nodes and amount of memory required. There are also tools available for verifying `makeflow` recipe files, as well as simple progress monitoring tools. Taken together, this package represents a powerful method for overseeing the execution of a given workflow, and serves as the backing pipeline management tool for the `hera_opm` package.

Listing 1 shows a sample TOML file where a user defines a workflow. The first section, called “Options”, specifies high-level options that are used by the `hera_opm` package in the construction of the executable script files. For instance, it allows the user to provide a default number of processors or memory quantity for batch jobs, or an Anaconda environment that should be activated before the execution of the work script. These options are fully enumerated in the documentation of the `hera_opm` package. Following the Options section, the “WorkFlow” section provides the main execution order for the pipeline. Further below, options are specified for each step, including the arguments that are passed in to the execution script. A corresponding batch job is generated for each step for each specified input file, with appropriate substitutions made for the file name and any necessary command line arguments. These batch job files are relatively simple, and only consist of: (i) creating an appropriate environment, (ii) calling the shell script which performs the required task, (iii) generating a “success” file if execution succeeded or a “failure” file if not. The `makeflow` program uses the existence of the “success” files to signify the task is completed and execution should continue on to the next job. If not, the task will be retried a certain number of times before being abandoned.

The shell scripts associated with each task follow a particular convention of being named for their corresponding step. For example, the script for running the `SETUP` task is called `do_SETUP.sh`. These scripts in turn can call scripts that perform the actual work associated with the task, such as Python scripts, C programs, or command-

²⁰https://github.com/HERA-Team/hera_mc

²¹<https://github.com/toml-lang/toml>

²²<https://github.com/cooperative-computing-lab/cctools>

```

[Options]
makeflow_type = "analysis"
path_to_do_scripts = "/path/to/task_scripts"
source_script = "~/.bashrc_hera"
conda_env = "hera"
base_mem = 8000
base_cpu = 1
timeout = "24h"

[ANT_METRICS_OPTS]
cross_cut = 5.0
dead_cut = 5.0
extension = ".ant_metrics.hdf5"

[XRFI_OPTS]
kt_size = 8
kf_size = 8

[WorkFlow]
actions = [
    "SETUP",
    "ANT_METRICS",
    "ADD_LIBRARIAN_ANT_METRICS",
    "XRFI",
    "ADD_LIBRARIAN_XRFI",
    "TEARDOWN"
]

[ANT_METRICS]
args = [
    "${basename}",
    "${ANT_METRICS_OPTS:cross_cut}",
    "${ANT_METRICS_OPTS:dead_cut}"
]

[ADD_LIBRARIAN_ANT_METRICS]
args = [
    "${basename}",
    "${ANT_METRICS_OPTS:extension}"
]

prereqs = "ANT_METRICS"

[XRFI]
queue = "gpu"
chunk_size = 10
stride_length = 10
time_centered = true
collect_stragglers = true
args = [
    "${basename}",
    "${XRFI_OPTS:kt_size}",
    "${XRFI_OPTS:kf_size}"
]

[ADD_LIBRARIAN_XRFI]
args = ["${basename}"]
prereqs = "XRFI"

```

Listing 1: A sample configuration file for the `hera.opm` package.

line utilities. The path to these scripts is specified in the configuration file, and the existence of these scripts is treated as a pre-requisite for the step, as a means of ensuring they exist before the pipeline is started. In the specific use case of invoking Python scripts and libraries, the `hera.opm` infrastructure allows for activating a target Anaconda environment, which can be specified as part of the configuration file. Besides this feature, the software environment is managed by the user and inherited from the execution shell.

The input for generating a pipeline are a configuration file and a list of data files to be processed. Based on the contents of the configuration file and the data files, the `hera.opm` package will generate a workflow script that can be parsed by `makeflow` and a series of shell scripts that will be run. In general, there is one shell script per step, per file, though there is support for processing “chunks” of input files as part of a single job. Once the workflow and shell scripts are generated, `makeflow` is used to manage the execution of the shell scripts. Successful completion of a shell script generates a placeholder output file that signifies successful completion, at which point `makeflow` launches the next job in the workflow. Once all the jobs have been completed successfully (or alternatively, if a job has failed more than an allowed number of times), `makeflow` halts execution of the pipeline as “finished.”

Although there is no formal requirement on the amount of work allowed as an individual step in the pipeline, in general tasks are encouraged to be as atomic as possible. For example, one might want to generate RFI flags for each *observation*, harmonize them into a single “per-observation” mask, and then apply them to all raw data products. Rather than performing all such tasks as a single job, the throughput of the total workflow may be improved by dividing the work into three separate jobs. Such an approach helps take advantage of the parallelism that execution on multiple compute nodes (or multiple cores on a single node) can offer. It also helps identify portions of the pipeline prone to failure more easily, and potentially resume stalled jobs after fixing a bug in a specific step. The atomic approach can have drawbacks, though if there is significant re-use of data between steps. For steps that are dominated by I/O rather than compute capability, it may actually be beneficial to perform several closely related steps as a single operation in the pipeline rather than using an atomic approach. Another concern is that including too many atomic steps can generate a workflow graph that becomes too large to compute in practice, so a trade-off may be necessary for a large number of workflow steps and input files. In general, profiling of various pipeline steps is beneficial. We discuss some performance profiling tools used as part of RTP below in Sec. 3.4.

One feature worth mentioning in the description of the workflow are the “setup” and “teardown” steps. These are operations that are meant to be performed once at the beginning and end of the workflow run, respectively. For instance, the setup step may pull and install the latest




Figure 1: A schematic diagram showing the various components of computing hardware in the on-site HERA cluster, and how the Librarian and RTP systems are distributed across them. In addition to the on-site cluster, an off-site processing facility is shown, which represents the cluster at the National Radio Astronomy Observatory (NRAO). Off-site processing at the Ilifu research cloud infrastructure in South Africa is handled similarly. Transfers between Librarian instances is facilitated by Globus (<https://www.globus.org/>), a service that facilitates large-scale data transfer between computing facilities. See Sec. 3.3 for further discussion of the interconnection of multiple sites.

versions of reference data for software packages, and the teardown may remove all orphaned files not cleaned up as part of the workflow. Their presence in the configuration file is optional, though if they are present, the setup and teardown steps must be the first and the last ones, respectively. As with other steps in the workflow, it is possible to specify arguments for them.

3.2. On-Site Analysis Workflows

The on-site data analysis for HERA has the constraint of looking at data in “real time”, i.e., shortly after they are taken. In addition to this time constraint, the data volume must also be reduced by a factor of 100 or more in order to move it within the network bandwidth available. We do not explore in detail the analysis steps taken, but instead focus on some of the requirements of various steps to highlight several features of the `hera_opm` package. The generation of a pipeline workflow file and shell scripts

is performed automatically by a monitoring system, after which the execution of the pipeline begins.

As mentioned above, the `makeflow` program is used to control the overall flow of the pipeline steps. One important feature of `makeflow` is its ability to interface with several different cluster resource managers. For the on-site pipeline in HERA, we make use of Slurm, which helps distribute the work of various tasks across a heterogeneous series of compute nodes. In particular, several analysis steps can take advantage of graphics processing units (GPUs) to accelerate calculations, and assigning these nodes to separate partitions in the Slurm cluster allows for specific tasks to be assigned to them. These computing options are specified in the configuration script (as for the XRFI step seen in Listing 1), and are handled by Slurm.

Due to the processing constraints imposed by the data rate, the correlator writes files that are only 16 seconds in length. For some steps, such as the RFI flagging done in the XRFI step, analyzing many contiguous time samples at once is scientifically important. In the case of RFI, multiple time samples help define a “baseline” for the data, so that flagging anomalous time-variable behavior is statistically simpler. For situations such as these, the RTP system allows the user to specify options for handling an automatic partitioning of the data into time-contiguous chunks. These options include: the number of files to include in a single time-contiguous chunk (up to and including all files in the workflow); the stride between these chunks (which in general can be different than the number of files in a chunk, if for example one wants to run analysis on overlapping chunks of data); whether the chunks are symmetric about a central file when generating groups, or are simply counted off; and how to handle an incomplete final group of files that is not evenly partitioned by the group size. The workflow designed by `hera_opm` accordingly tracks the files that belong to each group individually, so that downstream pre-requisites are handled automatically and correctly.

Figure 1 shows a schematic diagram of how the on-site cluster is organized, including the machines that run RTP and the Librarian system (discussed further below in Sec. 4). The RTP system is primarily concerned with running analysis, and makes use of the compute nodes. These nodes are assigned to various partitions in the Slurm system, and are available to run batch computation jobs. The head node serves as the submission host for Slurm jobs, as well as generating the initial workflow job using `hera_opm` when new data have been generated. The main raw data storage nodes are mounted on both the head node and each of the compute nodes via NFS mount (shown as a dashed line in the diagram), which allows for sharing data amongst the various nodes without the need to duplicate data. Intermediate data products needed for further analysis are written to the raw storage space. Output products earmarked for permanent storage and transfer back to US clusters are uploaded to the Librarian. In general, tasks which upload files to the Librarian are separate jobs in the

workflow. As mentioned above in Sec. 3.1, the steps in the workflow are made as atomic as possible. For example, in the sample workflow shown in Listing 1, steps uploading data to the Librarian are included in the workflow as separate tasks, as are analysis steps which feature different access patterns through the data.

3.3. Off-Site Post-Processing

An important feature of the `hera_opm` module is its portability: in principle, it is possible to run in any cluster environment that uses cluster resource management systems, or even on personal laptops. Because the machinery of `hera_opm` relies only on the Python-based package and `makeflow` for handling the computational flow, virtually any system should be able to make use of the infrastructure. In particular, HERA has access to a computing facility hosted by the National Radio Astronomy Observatory (NRAO). This cluster has a very different set of computing resources available compared to the on-site cluster, but the process of generating a workflow based on inputs using `hera_opm` is identical. The NRAO cluster uses the TORQUE resource manager, which is supported as a `makeflow` batch computing system. In processing the computing options in the configuration file, `hera_opm` is aware of differences in options between the two systems, and so the user is not required to make any changes in the configuration file apart from specifying a different queue system. When running the actual workflow, these steps are executed in a similar fashion. Assuming the input files and installed versions of software libraries are the same, then the output will be identical. The support for running on multiple clusters easily allows for portability of the entire analysis stack, and supports straightforward verification of any analysis output products.

The portability of the `hera_opm` package also allows for breaking up portions of the RTP system across different machines. Although currently all of the processing is run as part of the on-site cluster, the ability to easily and reliably make use of multiple computing environments opens the possibility of using multiple clusters for different portions of the RTP tasks. For instance, there are significant computing and processing resources available as part of the Ilifu research cloud infrastructure²³ of South Africa, which supports MeerKAT and SKA Pathfinder telescopes in their processing needs. Although additional infrastructure is required to ensure coordinated execution on various remote hosts, distributing the overall resource requirements across clusters is an overall beneficial approach to ensuring the required computing is accomplished in the requisite time window.

3.4. Pipeline Monitoring

As discussed in Sec. 3.1, RTP as a system is primarily concerned with executing a series of jobs in a user-defined workflow. However, there are several monitoring

tools built into RTP that help remote observers monitor the overall state and health of the system. In this section, we briefly cover some of the tools developed to make this monitoring possible. We make use of PostgreSQL²⁴ (PSQL), InfluxDB,²⁵ Chronograf,²⁶ and Grafana.²⁷ These subsystems allow for externally monitoring pipeline status without the need to directly log into the on-site system, and provides users with time-series data regarding key metrics related to the system. With the use of Chronograf and Grafana, metrics and statistics are plotted in real time without need for direct on-site access.

The key metrics, such as information about the state of processing systems, are stored in PSQL and InfluxDB servers running on-site. The database is populated by a system agent running on each computation host. This agent provides information such as the computational load, memory usage, and disk capacity as a function of time which are stored in the InfluxDB. It also contains key metrics related to the state of the correlator which are sent to the PSQL database. This information on the health of the processing system is also tracked by the Monitoring & Control subsystem, though a full description of this system is outside the scope of the current discussion. As it pertains to the status of the RTP system, these metrics are important for making sure the system is operating as intended, and has not stalled. The full data rate of 50 TB per night necessitates decreasing the total volume of data that are stored indefinitely, and so processing must continue with high reliability. Having tools for assessing and measuring these processes is essential to ensuring that HERA can deliver on its science goals.

Figure 2 shows a sample of the HERA dashboard, which contains information about the status of the telescope itself and data processing. On the left, this particular page shows the autocorrelation spectra from active antennas in near-real time, which can be useful for identifying poorly performing components. The user can easily select particular antennas to show, or get information for specific sections of the array. The right part of the page shows the layout of the entire array, with individual antennas color-coded based on their performance. In this particular image, the large number of red antennas denote ones which have not yet been connected, as HERA is still partially under construction at time of writing. In addition to these high-level views of the system, additional pages display more detailed information about logs from specific machines, the historical status of individual antennas, the fraction of files that have been successfully processed by RTP and moved to the USA for further analysis, and more. Taken together, this monitoring system allows all members of the collaboration an easy and straightforward way to en-

²³<http://www.ilifu.ac.za/>

²⁴<https://www.postgresql.org/>

²⁵<https://www.influxdata.com/>

²⁶<https://www.influxdata.com/time-series-platform/chronograf/>

²⁷<https://grafana.com/>




Figure 2: An example of the HERA dashboard. The top panel shows the current weather for the HERA site in the Karoo Desert of South Africa. The lower left panel shows live autocorrelations for the telescope, and allows the user to select specific antennas or sections of the array to be displayed. The lower right panel shows the status of individual antennas as part of the HERA array. In addition to the panels shown here, there are several other pages that show the status of other components of the telescope. The website is accessible both on-site and off-site which allows remote observers to understand the current state of the system.

sure the array is behaving as expected, which increases the amount of science-quality data generated by the telescope.

4. Long Term Storage and Data Transfer: The Librarian

We have also developed a new system which handles large data volumes and automated data processing specifically addressing several needs not covered by existing systems, as described in Sec. 2.1. The Librarian system addresses these requirements with a database backed program that records file locations as part of metadata, handles moves, and provides an API suitable for queries by automated pipelines or manual user control via a web page. As with the RTP system, it is written in Python, which allows for easily installing on a wide variety of systems. Although execution of commands on remote hosts assumes a Linux-based environment, these portions of the code could be easily replaced with analogous commands for other systems.

The HERA RTP system interacts closely with the Librarian, which provides long-term storage and data management capabilities for the HERA collaboration. At scale, HERA is projected to write 50 TB of data spread across 50,000 unique files each night. The Librarian system is designed to keep pace with the required data ingest operations, provide an interface for access to the accumulated HERA data set, and implement data transfer allowing members of the international HERA collaboration to perform a wide range of analyses at their local institutions.

4.1. Architecture

The Librarian system is composed of a loose federation of independently-operating *sites*. Each site hosts a freestanding Librarian server, which manages a database of file information, and *storage nodes* where file contents are actually stored. Each server exposes a JSON-based HTTP API that can be accessed by *clients* using several means: a Python library (the `hera-librarian` package), a command-line client, or an HTML web interface. Clients may be humans, RTP tasks and other local automated systems, or remote services such as other Librarian servers performing site-to-site synchronizations.

A key element of the Librarian design is its multi-site architecture. In the HERA collaboration, the two main Librarian sites are the observatory itself, where raw data are generated, and NRAO Domenici Science Operations Center (DSOC) in Socorro, NM, USA, which provides computing support to the collaboration. Each site runs its own Librarian server and mostly operates independently. At the same time, different sites can communicate and send data to one another as needed. In routine operations, transfers from South Africa to the USA happen on a daily basis as new data are recorded. However, if anything prevents site-to-site communication, such as loss of internet connectivity at the Karoo, users experience no issues besides a lack of the newest data. In terms of the CAP theorem (Gilbert and Lynch, 2002), which refers to the choice between *consistency* or *availability* in the presence of *partition tolerance* (i.e., network transmission failures), the Librarian prioritizes availability far ahead of consistency.

In typical cases, data flow through the HERA Librarian system in a radial pattern: raw data are generated in

South Africa and transferred “outward” to the USA and other processing sites. Despite this, the Librarian system is designed to support more varied flow topologies. For instance, an analysis of raw data might be run in the USA to generate new calibration files, which are then synchronized to South Africa to be used in RTP processing, creating a circular data flow.

The existence of multiple copies throughout the global system of Librarian instances poses a variant of the problem solved by version trackers like `git` or `mercurial` which we address with a similar solution to the one used by `git`: a two-tiered, append-only data model. This model has three main components:

1. The Librarian distinguishes between *files* and *file instances*. In the Librarian, a “file” is defined by its metadata: name, size, MD5 digest, creation time, and so on. A “file instance” is an actual copy of the file on a storage node. A Librarian server may hold zero, one, or many instances of a file.
2. All file metadata and file instance data are immutable.
3. File names are globally unique. Files can therefore be uniquely identified by their names.

Each Librarian server’s database of files may therefore adopt append-only semantics. If files are to be sent from one Librarian site to another, the metadata are easy and inexpensive to synchronize. Actual file instances are more expensive to synchronize, because they are generally much larger, but their immutability means that such synchronization can always eventually succeed.

In this model, the convention for choosing file names is essential to the system’s stability. While a single Librarian server can refuse to “mint” a new file record with a name that it knows to have already been taken, if two different servers mint files with the same name, there is presently no mechanism to reconcile their disagreement. In the experience of the HERA collaboration, file name clashes have never been a problem because the HERA Librarian sites generate different categories of data files (raw data, calibrations, data releases) and each such category has a distinctive file naming scheme. This convention maps well to real-world operations of scientific collaborations.

A “file” in the Librarian does not need to be a single traditional file. The Librarian software can handle directory trees as single coherent “files” for its indexing purposes, adopting a simple prescription for generating a reproducible MD5 digest. Though originally born out of necessity because several radio astronomy data formats are directory-based, this behaviour is useful for combining multiple files into a single unit before uploading to the Librarian, without the need to join the files together using `tar` or a similar command line utility.

Finally, each Librarian server also maintains a loosely-structured log of *file events*. These record simple operational data (e.g., “a new instance of this file was created on storage node X”) and housekeeping information such

as records of successful “standing order” transfers as described below. File events are not synchronized between sites.

The Librarian server uses the Tornado²⁸ framework for HTTP services and keeps all information in a PSQL relational database. It is statically configured with knowledge of local storage nodes, which are typically separate machines equipped with RAID arrays or other large-format devices. Storage nodes must be accessible over `ssh`, and must have some small Librarian utility programs installed, but do not need to run any special server software. Data transfers in and out of the Librarian system occur directly between clients and storage nodes: the server determines the commands necessary to move files between stores and executes them via remote execution on stores, rather than launching the transfers directly on the server.

4.2. Data Ingest

Data ingest into the Librarian is performed in a two-step process to ensure reliable operation. In the first phase, the client computes key metadata such as file size and MD5 digest and reports them to the server. If the upload is allowed, the server creates space in a “staging area” on one of the storage nodes and instructs the client to upload its data. In the second phase, the client uploads the data and, if the upload completes successfully, notifies the server of that fact. The server (via support tools installed on the storage node) verifies that the uploaded data agree with the client’s claims and, if so, atomically adds them to its data holdings. Because the Librarian has made its own copy of the file, the user is free to choose whether to delete the original copy upon the successful completion of an upload operation or retain a local copy.

While some file metadata are fully generic (e.g., size), the Librarian tracks additional data that are more specific to its use as an astronomical data manager. Of particular note is that each file is associated with an *observation* record defined by a start time (expressed as a Julian date, JD) and a duration. By convention, the JD is saved in the filename of raw HERA data, so derived data products which share a common prefix with these raw data can be matched to the same observation. When a derived data product is created from inputs corresponding to more than one *observation*, it is generally named according to the earliest *observation* of all of its inputs. While simple, this grouping scheme has worked well for the HERA collaboration’s usage patterns.

The Librarian has special support for one further layer of grouping. The user may instruct the Librarian server to group observations into *observing sessions*, each session ideally corresponding to a contiguous block of *observations* spanning a full night of observing. This assignment is done automatically and progressively: when so instructed, the

²⁸<https://github.com/tornadoweb/tornado>

Most Recent Sessions

Session ID	Date (UTC)	Start JD	Stop JD	Duration (days)	Num. Obs.	Num. Files (with instances)	Total size
1305915071	2021-05-24	2459359.25756645	2459359.25833024	0.0007637897506356239	2	10 (10)	195.9 MB
1305833276	2021-05-23	2459358.31086284	2459358.33556212	0.02469928003847599	109	653 (653)	10.8 GB
1305828657	2021-05-23	2459358.25739945	2459358.31117923	0.053779779933393	239	1483 (1483)	23.8 GB
1305742203	2021-05-22	2459357.25678191	2459357.25911157	0.0023296601139009	7	29 (29)	1.2 GB
1305655808	2021-05-21	2459356.25683394	2459356.34036532	0.08353138016536832	372	2432 (2432)	34.9 GB
1305569451	2021-05-20	2459355.25733768	2459355.34086907	0.08353138994425535	372	2273 (2273)	15.3 GB
1305483017	2021-05-19	2459354.25693944	2459354.34047082	0.08353138016536832	372	2417 (2417)	63.6 GB

Most Recent Files

Name	Created	Observation	Type	Source	Size
maint.20210706.nmasc.librarian.dbbbackup.pgdump details	2021-07-06 18:00:48	null	pgdump	local-maint	131.7 MB
maint.20210706.karoo.mandc.dbbbackup.pgdump details	2021-07-06 05:36:16	null	pgdump	karoo	7.9 GB
maint.20210706.karoo.librarian.dbbbackup.pgdump details	2021-07-06 05:24:53	null	pgdump	karoo	432.2 MB
maint.20210629.nmasc.librarian.dbbbackup.pgdump details	2021-06-29 18:00:47	null	pgdump	local-maint	131.7 MB
maint.20210629.karoo.mandc.dbbbackup.pgdump details	2021-06-29 05:36:03	null	pgdump	karoo	7.9 GB

Figure 3: A sample of the Librarian web interface. A list of recent observation sessions is at the top, with more detailed information below. Clicking on the number identifying an observation session takes the user to a page that provides additional details about the data, with similar behavior for the identifying numbers of individual files. From here users can download individual files or (more commonly) make a copy of datasets on fast cluster computing storage from the slower main Librarian storage.

server determines sessions from the start times and durations of all *observations* that have not already been assigned to a session. Typically, this grouping is performed automatically at the end of uploading a night’s worth of observations to the Librarian. These observing session groupings are then used by the RTP system for defining a workflow and performing analysis. Observation and observing session parameters are immutable once assigned.

4.3. Searches

The Librarian is equipped with a simple search framework. Besides allowing human users to explore the Librarian data holdings, the search infrastructure is also intended to be a useful building block for system automation tasks, such as the “standing order” system which moves data between sites. Standing orders are defined in terms of search queries and evaluated periodically (typically a few times per hour). Files that match the search criteria for standing orders are added to a list of transfers to be performed. This mechanism allows for automatic data transfer between, e.g., the on-site Librarian and the one at NRAO, as shown diagrammatically in Figure 1.

Searches in the Librarian are expressed as JSON documents that are transcribed into SQL queries. Searches

are represented in the JSON document as trees of clauses composed using standard boolean operations. Most search clauses map directly into SQL `SELECT` terms: for instance, a “name-like” JSON term is translated by the Librarian’s search API into a `name LIKE` query in the resulting SQL. Searches may query the Librarian’s tables of files, observations, or observing sessions. Search results for file queries can yield structured file metadata, lists of file names, or lists of file instances. For HERA, these metadata include the time and duration of individual observations, which “observing session” an individual observation belongs to, and the local sidereal time of an observation.

Certain search clauses can be fairly expensive to evaluate. For instance, one can search on the number of files associated with a given observing session, which is useful for automating checks for serious failures in the data recording system. The Librarian computes this number on-the-fly every time that it is needed, performing a query on its backing database that must join over the tables of files, observations, and observing sessions.

4.4. Site-To-Site Data Transfers

Data transfers between Librarian sites follow the same two-phase process as is used for intra-site, client-to-server,

transfers. Each Librarian server may be configured to have “standing orders” to automate data transfers. Each standing order is defined by a file search query and a destination site. The server periodically evaluates the query and identifies any files that have not yet been sent to the destination. Transfers of these files are then initiated. These uploads are performed with modest parallelization because multiple streams can make use of the available network bandwidth more effectively than a single one. When each file has been successfully transferred, a record of the transfer is made in its event log, preventing future upload attempts. As with all aspects of the Librarian, this design of this feature is intended to provide reliable operation in the face of unreliable network availability between sites.

Inter-Librarian data transfers can make use of Globus, a service that facilitates large-scale data transfer between computing facilities. Globus primarily serves facilities that support scientific research, though it also supports transfer to commercial computing infrastructure like Amazon Web Services (AWS²⁹). Many supercomputing facilities are already established as Globus endpoints, which allows users to tap into their infrastructure. In addition to tools that automatically adapt for various network configurations, Globus also provides detailed information on file transfer status, robust logs for records, and notifications on successful (or unsuccessful) transfers. Setup of a Personal Globus Endpoint on the Librarian server is straightforward, and only requires a few additional entries in the Librarian configuration file. In the case of HERA, the NRAO computing facility is a Globus endpoint, and the computing facility on-site makes use of Personal Globus Endpoints to initiate transfers to the NRAO facility. In this way, HERA data transfers are able to offload some of the logistical issues related to data transfer by using Globus.

4.5. Client Interfaces for Access to the Librarian

Each Librarian server supports client access through a command-line tool, a Python library, and an HTML web interface. Most client access is ultimately provided through the `hera_librarian` Python package, which has few dependencies and so can be easily installed by the end user. The user configures the client with knowledge of one or more Librarian servers through a small configuration file. The `hera_librarian` module then provides fairly direct access to the JSON APIs provided by the remote servers. For example, a script might implement a slightly more sophisticated version of the “standing order” system, query a local Librarian server for a particular series files, or apply a filter that is too complex to implement server-side. More concretely, the HERA operations team has used the Python API to query a remote Librarian to check which of the remaining files it holds, and then initiating a transfer from the local Librarian server to the remote one. The Python API provides the most powerful and flexible way

of interacting with a Librarian server, which makes it a valuable tool when using the Librarian in practice.

The `hera_librarian` Python package provides a command-line tool, `librarian`, that provides a more convenient interface to the most common Librarian access tasks, such as uploading new files to the librarian and searching for existing files. For this latter task of searching for files, only a few key bits of information are provided, such as the file name and size on disk. For more detailed information, the user must use the Python API or the web-based interface. Users can install the `hera_librarian` package in a local Python environment, and connect to either local or remote Librarian servers. The information details of accessing servers is handled via a per-user configuration file, which may require port forwarding via an SSH connection to access servers behind a remote firewall.

Finally, each Librarian server provides an HTML user interface along with its JSON API. While a rudimentary authentication scheme is implemented, the primary means of access control is SSH authentication. HTTP interfaces to the Librarian are not exposed on the open Internet: as with using the Python API, SSH port-forwarding may be required to access the web server.

Figure 3 shows a sample of the Librarian web interface. At the top is a summary of the most recent Observing Sessions, including their Julian Date and the number of constituent Observations. Clicking on the Observing Session identification number will take the user to a page with more detailed information. Below the section of Observing Sessions, the Librarian displays the most recent files added to it. This list merely provides a quick overview of the most recent files, and is not intended to be an exhaustive list of all files in the Librarian. The web interface also has a search page, where the user may construct a query matching specific files more easily. From here users can download individual files or (more commonly) make a copy of datasets on fast cluster computing storage from the slower main Librarian storage.

The Librarian web interface is intentionally basic, because the expectation and intention is that day-to-day data analysis should be driven through command-line access or Python scripts while on site the system is tracked by the HERA monitor and control system. While the web interface is helpful for administration and monitoring, with HERA’s large data volumes and extremely uniform data holdings, data retrieval should generally be done in some kind of automated fashion, at the beginning of an analysis pipeline.

There is no particular requirement that Librarian client libraries must be implemented in Python. The Librarian server APIs use standard HTTP and JSON patterns, and so it would be straightforward to implement client software in nearly any modern programming language.

4.6. Docker Container

In addition to the “bare-metal” installation of Librarian described above, the `librarian` package also supports in-

²⁹<https://aws.amazon.com>

stallation as a series of Docker³⁰ containers. Rather than using a single monolithic container for the entire application, the different components are separated into different images. This takes advantage of orchestration software for container management to facilitate an overall higher up-time and reliability. The Librarian package is composed of three containers: (i) the backing PostgreSQL database, (ii) the Python-based Librarian application, and (iii) a “store” server for saving data files. The latter two containers are based off the same base image, which includes an installation of the `hera-librarian` package. The “store” image also runs an ssh server, so that minimal changes are required to the traditional method of operation of the Librarian. The networking interfaces between these containers is specified in a `docker-compose.yml` file. A full series of instructions is available in the repository describing the installation procedure. Note that the SQL database and the location of the store server where data are saved should be configured as Volumes, so that the data and metadata persist between Docker image instances.

5. Summary

The HERA project has built on lessons from several projects and through a series of iterations arrived at a system for processing and managing large data volumes. Formative events in prior projects included loss of half a season of data to accidental deletion, movement of data using frequent shipments, inadvertent siloing of data inside an unscriptable archive system, and failure of automated scripts leading to silent telescope shutdowns. These problems largely stemmed from a paucity of automation and are largely in the past. The system described here is currently in operation at several sites around the world processing 10 TB per day at peak operation—a rate which is expected to grow to hundreds of TB per day in the coming months.

With the explosion of raw data that has come about in astronomy research, data analysis and processing has become equally important. Developing flexible yet reliable frameworks for analyzing and storing data are of the utmost importance for providing trustworthy scientific results. We have presented the RTP⁵ and the Librarian⁶ systems as they pertain to data analysis in HERA, which has particularly strict requirements for processing and data handling. However, these tools have been developed such that they are modular and general enough to be applied in other astronomy contexts, such as in experiments outside of radio astronomy or future projects like the SKA. The source code underlying these systems is free and open source, and available on GitHub. We hope that these tools may be useful to other collaborations or projects as the volume of astronomy data will only increase in the future.

In the future, we plan to increase the amount of automation and transparency that happen in the system. In particular, visualizing the status of the RTP system through the monitoring website shown in Figure 2 will be essential for ensuring that data are successfully processed in a timely fashion. We also plan to develop the multi-site processing paradigm discussed in Sec. 3.3 further, as the relatively limited computational resources on site necessitate using additional compute resources. Throughout all of these development efforts, the ability to test and deploy the code both locally and at scale has been facilitated by continuous integration testing principles (e.g., writing a robust series of tests which must pass before changes can be committed to the main software repo), leading to higher quality code and better reliability. Whether adapting these RTP and Librarian systems or building bespoke ones, users facing similar design considerations are encouraged to use robust testing frameworks for ensuring reliable operation of critical software infrastructure components.

Acknowledgment

This material is based upon work supported by the National Science Foundation under Grant No. 1636646, the Gordon and Betty Moore Foundation through grant GBMF5215 to the Massachusetts Institute of Technology, and institutional support from the HERA collaboration partners. HERA is hosted by the South African Radio Astronomy Observatory, which is a facility of the National Research Foundation, an agency of the Department of Science and Innovation. 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. G. Bernardi acknowledges funding from the INAF PRIN-SKA 2017 project 1.05.01.88.04 (FORECaST), support from the Ministero degli Affari Esteri della Cooperazione Internazionale – Direzione Generale per la Promozione del Sistema Paese Progetto di Grande Rilevanza ZA18GR02 and the National Research Foundation of South Africa (Grant Number 113121) as part of the ISARP RADIOSKY2020 Joint Research Scheme, from the Royal Society and the Newton Fund under grant NA150184 and from the National Research Foundation of South Africa (grant No. 103424). P. Bull acknowledges funding for part of this research from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant agreement No. 948764), and from STFC Grant ST/T000341/1. J. S. Dillon gratefully acknowledges the support of the NSF AAPF award #1701536. N Kern acknowledges support from the MIT Pappalardo fellowship. A. Liu acknowledges support from the New Frontiers in Research Fund Exploration grant program, the Canadian Institute for Advanced Research (CIFAR) Azrieli Global Scholars program, a Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant and a Discovery Launch Supplement, the

³⁰<https://www.docker.com/>

Sloan Research Fellowship, and the William Dawson Scholarship at McGill.

References

- Aguirre, J.E., Murray, S.G., Pascua, R., Martinot, Z.E., Burba, J., Dillon, J.S., Jacobs, D.C., Kern, N.S., Kittiwisit, P., Kolopanis, M., Lanman, A., Liu, A., Whitler, L., Abdurashidova, Z., Alexander, P., Ali, Z.S., Balfour, Y., Beardsley, A.P., Bernardi, G., Billings, T.S., Bowman, J.D., Bradley, R.F., Bull, P., Carey, S., Carilli, C.L., Cheng, C., DeBoer, D.R., Dexter, M., de Lera Acedo, E., Ely, J., Ewall-Wice, A., Fagnoni, N., Fritz, R., Furlanetto, S.R., Gale-Sides, K., Glendenning, B., Gorthi, D., Greig, B., Grobbelaar, J., Halday, Z., Hazelton, B.J., Hewitt, J.N., Hickish, J., Julius, A., Kerrigan, J., Kohn, S.A., La Plante, P., Lekalake, T., Lewis, D., MacMahon, D., Malan, L., Malgas, C., Maree, M., Matsetela, E., Mesinger, A., Molewa, M., Morales, M.F., Mosiane, T., Neben, A.R., Bikolic, B., Parsons, A.R., Patra, N., Pieterse, S., Pober, J.C., Razavi-Ghods, N., Ringuette, J., Robnett, J., Rosie, K., Santos, M., Sims, P., Smith, C., Syce, A., Thyagarajan, N., Williams, P.K.G., Zheng, H., 2021. Validation of the HERA Phase I Epoch of Reionization 21 cm Power Spectrum Software Pipeline. in prep. .
- Albrecht, J., Alves, Antonio Augusto, J., Amadio, G., Andronico, G., Anh-Ky, N., Aphaecetche, L., Apostolakis, J., Asai, M., Atzori, L., Babik, M., Bagliesi, G., Bandieramonte, M., Banerjee, S., Barisits, M., Bauerdick, L.A.T., Belforte, S., Benjamin, D., Bernius, C., Bhimji, W., Bianchi, R.M., Bird, I., Biscarat, C., Blomer, J., Bloom, K., Boccali, T., Bockelman, B., Bold, T., Bonacorsi, D., Boveia, A., Bozzi, C., Bracko, M., Britton, D., Buckley, A., Buncic, P., Calafiura, P., Campana, S., Canal, P., Canali, L., Carlino, G., Castro, N., Cattaneo, M., Cerminara, G., Cervantes Villanueva, J., Chang, P., Chapman, J., Chen, G., Childers, T., Clarke, P., Clemencic, M., Cogneras, E., Coles, J., Collier, I., Colling, D., Corti, G., Cosmo, G., Costanzo, D., Couturier, B., Cranmer, K., Cranshaw, J., Cristella, L., Crooks, D., Crépé-Renaudin, S., Currie, R., Dallmeier-Tiessen, S., De, K., De Cian, M., De Roeck, A., Delgado Peris, A., Derue, F., Di Girolamo, A., Di Guida, S., Dimitrov, G., Doglioni, C., Dotti, A., Duellmann, D., Duflo, L., Dykstra, D., Dziejniewicz-Wojcik, K., Dziurda, A., Egede, U., Elmer, P., Elmsheuser, J., Elvira, V.D., Eulisse, G., Farrell, S., Ferber, T., Filipic, A., Fisk, I., Fitzpatrick, C., Flix, J., Formica, A., Forti, A., Franzoni, G., Frost, J., Fuess, S., Gaede, F., Ganis, G., Gardner, R., Garonne, V., Gellrich, A., Genser, K., George, S., Geurts, F., Gheata, A., Gheata, M., Giacomini, F., Giagu, S., Giffels, M., Gingrich, D., Girone, M., Gligorov, V.V., Glushkov, I., Gohn, W., Benito Gonzalez Lopez, J., González Caballero, I., González Fernández, J.R., Govi, G., Grandi, C., Grasland, H., Gray, H., Grillo, L., Guan, W., Gutsche, O., Gyurjyan, V., Hanushevsky, A., Hariri, F., Hartmann, T., Harvey, J., Hauth, T., Hegner, B., Heinemann, B., Heinrich, L., Heiss, A., Hernández, J.M., Hildreth, M., Hodgkinson, M., Hoeche, S., Holzman, B., Hristov, P., Huang, X., Ivanchenko, V.N., Ivanov, T., Iven, J., Jashal, B., Jayatilaka, B., Jones, R., Jouvin, M., Jun, S.Y., Kagan, M., Kalderon, C.W., Kane, M., Karavakis, E., Katz, D.S., Kcira, D., Keeble, O., Kersevan, B.P., Kirby, M., Klimentov, A., Klute, M., Komarov, I., Konstantinov, D., Koppenburg, P., Kowalkowski, J., Kreczko, L., Kuhr, T., Kutschke, R., Kuznetsov, V., Lampl, W., Lancon, E., Lange, D., Lassnig, M., Laycock, P., Leggett, C., Letts, J., Lewendel, B., Li, T., Lima, G., Linacre, J., Linden, T., Livny, M., Lo Presti, G., Lopienski, S., Love, P., Lyon, A., Magini, N., Marshall, Z.L., Martelli, E., Martin-Haugh, S., Mato, P., Mazumdar, K., McCauley, T., McFayden, J., McKee, S., McNab, A., Mehdiyev, R., Meinhard, H., Menasce, D., Mendez Lorenzo, P., Serhan Mete, A., Michelotto, M., Mitrevski, J., Moneta, L., Morgan, B., Mount, R., Moysse, E., Murray, S., Nairz, A., Neubauer, M.S., Norman, A., Novaes, S., Novak, M., Oyanguren, A., Ozturk, N., Pacheco Pages, A., Paganini, M., Pansanel, J., Pascuzzi, V.R., Patrick, G., Pearce, A., Pearson, B., Pedro, K., Perdue, G., Perez-Calero Yzquierdo, A., Perrozzi, L., Petersen, T., Petric, M., Petzold, A., Piedra, J., Piilonen, L., Piparo, D., Pivarski, J., Pokorski, W., Polci, F., Potamianos, K., Psihas, F., Navarro, A.P., Quast, G., Raven, G., Reuter, J., Ribon, A., Rinaldi, L., Ritter, M., Robinson, J., Rodrigues, E., Roiser, S., Rousseau, D., Roy, G., Rybkine, G., Sailer, A., Sakuma, T., Santana, R., Sartirana, A., Schellman, H., Schovancová, J., Schramm, S., Schulz, M., Sciabà, A., Seidel, S., Sekmen, S., Serfon, C., Severini, H., Sexton-Kennedy, E., Seymour, M., Sgalaberna, D., Shapoval, I., Shiers, J., Shiu, J.G., Short, H., Piero Siroli, G., Skipsey, S., Smith, T., Snyder, S., Sokoloff, M.D., Spentzouris, P., Stadie, H., Stark, G., Stewart, G., Stewart, G.A., Sánchez, A., Sánchez-Hernández, A., Taffard, A., Tamponi, U., Templon, J., Tenaglia, G., Tsulaia, V., Tunnell, C., Vaandering, E., Valassi, A., Vallecorsa, S., Valsan, L., Van Gemmeren, P., Vernet, R., Viren, B., Vlimant, J.R., Voss, C., Votava, M., Vuosalo, C., Vázquez Sierra, C., Wartel, R., Watts, G.T., Wenaus, T., Wenzel, S., Williams, M., Winklmeier, F., Wissing, C., Wuethwein, F., Wynne, B., Xiaomei, Z., Yang, W., Yazgan, E., 2017. A Roadmap for HEP Software and Computing R&D for the 2020s. arXiv e-prints , arXiv:1712.06982arXiv:1712.06982.
- Albrecht, M., Donnelly, P., Bui, P., Thain, D., 2012. Makeflow: A portable abstraction for data intensive computing on clusters, clouds, and grids, in: Proceedings of the 1st ACM SIGMOD Workshop on Scalable Workflow Execution Engines and Technologies, Association for Computing Machinery, New York, NY, USA. URL: <https://doi.org/10.1145/2443416.2443417>, doi:10.1145/2443416.2443417.
- Carmichael, R., Braga-Henebry, P., Thain, D., Emrich, S., 2010. Biocompute: Towards a collaborative workspace for data intensive bio-science, in: Proceedings of the 19th ACM International Symposium on High Performance Distributed Computing, Association for Computing Machinery, New York, NY, USA. p. 489–498. URL: <https://doi.org/10.1145/1851476.1851547>, doi:10.1145/1851476.1851547.
- DeBoer, D.R., Parsons, A.R., Aguirre, J.E., Alexander, P., Ali, Z.S., Beardsley, A.P., Bernardi, G., Bowman, J.D., Bradley, R.F., Carilli, C.L., Cheng, C., de Lera Acedo, E., Dillon, J.S., Ewall-Wice, A., Fadana, G., Fagnoni, N., Fritz, R., Furlanetto, S.R., Glendenning, B., Greig, B., Grobbelaar, J., Hazelton, B.J., Hewitt, J.N., Hickish, J., Jacobs, D.C., Julius, A., Kariseb, M., Kohn, S.A., Lekalake, T., Liu, A., Loo, A., MacMahon, D., Malan, L., Malgas, C., Maree, M., Martinot, Z., Mathison, N., Matsetela, E., Mesinger, A., Morales, M.F., Neben, A.R., Patra, N., Pieterse, S., Pober, J.C., Razavi-Ghods, N., Ringuette, J., Robnett, J., Rosie, K., Sell, R., Smith, C., Syce, A., Tegmark, M., Thyagarajan, N., Williams, P.K.G., Zheng, H., 2017. Hydrogen Epoch of Reionization Array (HERA). PASP 129, 045001. doi:10.1088/1538-3873/129/974/045001, arXiv:1606.07473.
- Dillon, J.S., Lee, M., Ali, Z.S., Parsons, A.R., Orosz, N., Nunhokee, C.D., La Plante, P., Beardsley, A.P., Kern, N.S., Abdurashidova, Z., Aguirre, J.E., Alexander, P., Balfour, Y., Bernardi, G., Billings, T.S., Bowman, J.D., Bradley, R.F., Bull, P., Burba, J., Carey, S., Carilli, C.L., Cheng, C., DeBoer, D.R., Dexter, M., de Lera Acedo, E., Ely, J., Ewall-Wice, A., Fagnoni, N., Fritz, R., Furlanetto, S.R., Gale-Sides, K., Glendenning, B., Gorthi, D., Greig, B., Grobbelaar, J., Halday, Z., Hazelton, B.J., Hewitt, J.N., Hickish, J., Jacobs, D.C., Julius, A., Kerrigan, J., Kittiwisit, P., Kohn, S.A., Kolopanis, M., Lanman, A., Lekalake, T., Lewis, D., Liu, A., Ma, Y.Z., MacMahon, D., Malan, L., Malgas, C., Maree, M., Martinot, Z.E., Matsetela, E., Mesinger, A., Molewa, M., Morales, M.F., Mosiane, T., Murray, S., Neben, A.R., Nikolic, B., Pascua, R., Patra, N., Pieterse, S., Pober, J.C., Razavi-Ghods, N., Ringuette, J., Robnett, J., Rosie, K., Santos, M.G., Sims, P., Smith, C., Syce, A., Tegmark, M., Thyagarajan, N., Williams, P.K.G., Zheng, H., 2020. Redundant-baseline calibration of the hydrogen epoch of reionization array. Monthly Notices of the RAS 499, 5840–5861. doi:10.1093/mnras/staa3001, arXiv:2003.08399.
- Gilbert, S., Lynch, N., 2002. Brewer’s conjecture and the feasibility of consistent, available, partition-tolerant web services. SIGACT News 33, 51–59. URL: <https://doi.org/10.1145/>

- 564585.564601, doi:10.1145/564585.564601.
- Kern, N.S., Dillon, J.S., Parsons, A.R., Carilli, C.L., Bernardi, G., Abdurashidova, Z., Aguirre, J.E., Alexander, P., Ali, Z.S., Balfour, Y., Beardsley, A.P., Billings, T.S., Bowman, J.D., Bradley, R.F., Bull, P., Burba, J., Carey, S., Cheng, C., DeBoer, D.R., Dexter, M., de Lera Acedo, E., Ely, J., Ewall-Wice, A., Fagnoni, N., Fritz, R., Furlanetto, S.R., Gale-Sides, K., Glendenning, B., Gorthi, D., Greig, B., Grobbelaar, J., Halday, Z., Hazelton, B.J., Hewitt, J.N., Hickish, J., Jacobs, D.C., Julius, A., Kerrigan, J., Kittiwisit, P., Kohn, S.A., Kolopanis, M., Lanman, A., La Plante, P., Lekalake, T., Liu, A., MacMahon, D., Malan, L., Malgas, C., Maree, M., Martinot, Z.E., Matsetela, E., Mesinger, A., Molewa, M., Morales, M.F., Mosiane, T., Murray, S.G., Neben, A.R., Nikolic, B., Nunhokee, C.D., Patra, N., Pieterse, S., Pober, J.C., Razavi-Ghods, N., Ringuelette, J., Robnett, J., Rosie, K., Sims, P., Smith, C., Syce, A., Thyagarajan, N., Williams, P.K.G., Zheng, H., 2020a. Absolute Calibration Strategies for the Hydrogen Epoch of Reionization Array and Their Impact on the 21 cm Power Spectrum. *Astrophysical Journal* 890, 122. doi:10.3847/1538-4357/ab67bc, arXiv:1910.12943.
- Kern, N.S., Parsons, A.R., Dillon, J.S., Lanman, A.E., Liu, A., Bull, P., Ewall-Wice, A., Abdurashidova, Z., Aguirre, J.E., Alexander, P., Ali, Z.S., Balfour, Y., Beardsley, A.P., Bernardi, G., Bowman, J.D., Bradley, R.F., Burba, J., Carilli, C.L., Cheng, C., DeBoer, D.R., Dexter, M., de Lera Acedo, E., Fagnoni, N., Fritz, R., Furlanetto, S.R., Glendenning, B., Gorthi, D., Greig, B., Grobbelaar, J., Halday, Z., Hazelton, B.J., Hewitt, J.N., Hickish, J., Jacobs, D.C., Julius, A., Kerrigan, J., Kittiwisit, P., Kohn, S.A., Kolopanis, M., La Plante, P., Lekalake, T., MacMahon, D., Malan, L., Malgas, C., Maree, M., Martinot, Z.E., Matsetela, E., Mesinger, A., Molewa, M., Morales, M.F., Mosiane, T., Murray, S.G., Neben, A.R., Parsons, A.R., Patra, N., Pieterse, S., Pober, J.C., Razavi-Ghods, N., Ringuelette, J., Robnett, J., Rosie, K., Sims, P., Smith, C., Syce, A., Thyagarajan, N., Williams, P.K.G., Zheng, H., 2020b. Mitigating Internal Instrument Coupling for 21 cm Cosmology. II. A Method Demonstration with the Hydrogen Epoch of Reionization Array. *Astrophysical Journal* 888, 70. doi:10.3847/1538-4357/ab5e8a, arXiv:1909.11733.
- La Plante, P., Dillon, J.S., Kolopanis, M., Beardsley, A.P., Kern, N.S., Wilensky, M., HERA Collaboration, 2021. hera.opm: The HERA Online Processing Module. arXiv:2104.001.
- Pober, J.C., Parsons, A.R., Aguirre, J.E., Ali, Z., Bradley, R.F., Carilli, C.L., DeBoer, D., Dexter, M., Gugliucci, N.E., Jacobs, D.C., Klima, P.J., MacMahon, D., Manley, J., Moore, D.F., Stefan, I.I., Walbrugh, W.P., 2013. Opening the 21 cm Epoch of Reionization Window: Measurements of Foreground Isolation with PAPER. *Astrophysical Journal Letters* 768, L36. doi:10.1088/2041-8205/768/2/L36, arXiv:1301.7099.
- Spreeuw, H., Madougou, S., Van Haren, R., Weel, B., Belloum, A., Maassen, J., 2019. Unlocking the lofar lta, in: 2019 15th International Conference on eScience (eScience), pp. 467–470. doi:10.1109/eScience.2019.00061.
- The HERA Collaboration, Abdurashidova, Z., Aguirre, J.E., Alexander, P., Ali, Z.S., Balfour, Y., Beardsley, A.P., Bernardi, G., Billings, T.S., Bowman, J.D., Bradley, R.F., Bull, P., Burba, J., Carey, S., Carilli, C.L., Cheng, C., DeBoer, D.R., Dexter, M., de Lera Acedo, E., Dibblee-Barkman, T., Dillon, J.S., Ely, J., Ewall-Wice, A., Fagnoni, N., Fritz, R., Furlanetto, S.R., Gale-Sides, K., Glendenning, B., Gorthi, D., Greig, B., Grobbelaar, J., Halday, Z., Hazelton, B.J., Hewitt, J.N., Hickis, J., Jacobs, D.C., Julius, A., Kern, N.S., Kerrigan, J., Kittiwisit, P., Kohn, S.A., Kolopanis, M., Lanman, A., La Plante, P., Lekalake, T., Lewis, D., Liu, A., MacMahon, D., Malan, L., Malgas, C., Maree, M., Martinot, Z.E., Matsetela, E., Mesinger, A., Molewa, M., Morales, M.F., Mosiane, T., Murray, S.G., Neben, A.R., Nikolic, B., Nunhokee, C.D., Parsons, A.R., Patra, N., Pascua, R., Pieterse, S., Pober, J.C., Razavi-Ghods, N., Ringuelette, J., Robnett, J., Rosie, K., Sims, P., Singh, S., Smith, C., Syce, A., Thyagarajan, N., Williams, P.K.G., Zheng, H., 2021. First Results from HERA Phase I: Upper Limits on the Epoch of Reionization 21 cm Power Spectrum. in prep. .
- Wijnholds, S.J., Willis, A.G., Salvini, S., 2018. Baseline-dependent averaging in radio interferometry. *Monthly Notices of the RAS* 476, 2029–2039. doi:10.1093/mnras/sty360.
- Williams, P.K.G., La Plante, P., HERA Collaboration, 2021. Librarian: The HERA Librarian. arXiv:2104.002.