

Cover Page



Universiteit Leiden



The following handle holds various files of this Leiden University dissertation:
<http://hdl.handle.net/1887/81487>

Author: Mechev, A.P.

Title: Orchestration of Distributed LOFAR Workflows

Issue Date: 2019-12-09

As astronomical observatories produce ever-growing data-sets, the processing challenges for these data will continue to increase. Extensive astronomical surveys, expected to create petabytes of data, can no longer be processed on a single machine or small dedicated clusters at scientific institutions. Serving the scientific requirements of these surveys will require large scale distributed processing.

CERN's World-Wide computing Grid provides sufficient resources for such projects; however, its focus is on distributed Monte-Carlo simulations. This high-throughput infrastructure offers opportunities for parallel processing of radio astronomy data sets. To take advantage of these resources and implement complex astronomical workflows to a grid-like environment requires a framework to distribute and monitor jobs. Furthermore, processing and re-processing thousands of observations efficiently requires workflow orchestration software. We aim to enable the 30+ petabyte LOFAR Two-meter Sky Survey (LoTSS) by combining high throughput processing infrastructure with modern workflow orchestration software.

8.1 Summary of Thesis Contributions

The work in this thesis focuses on software built to accelerate, parallelize, and automate LOFAR processing as well as the insights obtained into large scale processing of LOFAR data. We have built a scalability model which we use to understand the performance of LOFAR broadband processing pipelines. Our model brings novel insights into the limits of our current pipelines, as well as suggestions to improve processing throughput.

We have built a platform for processing a radio astronomy data on a heterogeneous, distributed infrastructure. We exploit the data-level parallelism of com-

putationally intensive processing tasks, and our work makes several LOFAR scientific projects possible. Additionally, our insights into distributed execution of complex pipelines are crucial for enabling sizeable astronomical surveys. We expect distributed processing will become an increasingly important paradigm in astronomy.

Finally, we created an automated workflow system with the goal to automatically produce high fidelity images from LOFAR observations. This system was a successful integration of industry software into radio astronomy, one of the goals of this thesis and its associated grant. Bringing open-source tools used in industry is crucial to keeping long-lived scientific projects maintainable and productive. Our results were an important step in enabling high throughput, automated processing of LOFAR scientific workflows.

Our advances in understanding LOFAR processing inefficiencies, exploiting data-level parallelism, and automating workflows are important steps to modernizing LOFAR scientific processing. The lessons learned in this work can be directly applied in other scientific fields that need to process data at overwhelming rates.

8.2 Answers to Research Questions

Research Question 1: How can we use a distributed shared infrastructure for efficient LOFAR data processing?

In Chapters 2 and 3, we present our results enabling massively distributed processing of LOFAR data. We describe the underlying platform, inherited from the High Energy Physics community and the modifications to these tools that were required to host sophisticated processing software. We describe these modifications and discuss the resulting increase in throughput. Finally, we estimate the time saved by parallelizing LOFAR data processing. The work described in these chapters is essential to producing scientific data sets at a high rate, particularly considering the high data rates produced by LOFAR.

Research Question 2: How can we build software to effortlessly accelerate complex pipelines for radio astronomy?

Chapters 5 and 7 present our work on parallelizing LOFAR scientific pipelines on a distributed shared infrastructures. We integrate a mature workflow orchestration package with distributed LOFAR processing. We discuss the need for this orchestration, as well as the abilities to support additional complex pipelines.

As an example application, we build a Continuous Integration pipeline tasked with verifying and validating the initial steps of LOFAR processing.

***Research Question 3:** Can we automatically collect performance information during massively distributed processing and predict run times for future data sets?*

Chapters 4 and 6 describe a performance monitoring suite for LOFAR data and our scalability model for LOFAR processing. When running massively distributed processing, scientists are unable to monitor the performance of the underlying software. Collecting these statistics is necessary for understanding processing inefficiencies and suggest ways to accelerate data processing. Performance data can also be used to understand the effect of processing parameters on the resource usage of complex pipelines. We study this in detail, building a model that can be used to understand the scalability of multiple processing steps. This model shows the limitations imposed by available processing resources as well as suggestions on decreasing processing time without sacrificing scientific data quality.

8.3 Limitations

Using the software described, the LOFAR Surveys team was able to process several petabytes of archived data and produce scientific quality images. Despite the successes of the project, several issues occasionally impede data processing and prevent rapid deployment of software pipelines.

High throughput processing of LOFAR data requires the initial processing steps to be performed at the data archive locations. While deploying new versions of the LOFAR software and pipelines at SUFRsara is straightforward, the same is not true for other LTA locations. LOFAR data needs to be processed at LTA locations that do not support any modern containerization software nor other software distribution methods. This makes deploying new software is difficult and time-consuming. Additionally, orchestrating jobs at these sites requires additional integration with our job monitoring tools due to lack of internet access from some HTC clusters. Integrating locations not suited for large scale distributed processing is an ongoing challenge for the LoTSS survey and other LOFAR projects.

A further limitation is the authentication of our processing software and the authentication requirements for data access. Distributed processing of large data sets requires having access to the data at multiple locations, however the intermediate products produced by our workflows are not public and require an active x.509

certificate authorised by the LOFAR SKSP project. The current workaround to this limitation is to maintain active certificates at each processing location. Upcoming features of the dCache storage system, bearer tokens called macaroons, will make it possible to overcome this limitation.

Finally, our current software distribution does not assign long-term version numbers to software images and scripts, nor is there a way to store these images or cite them in related papers. Implementing proper versioning is crucial to not only making data processing easily reproducible, but also make it possible to recognize the effort put into building and distributing software images. Overcoming these limitations will enable FAIR science with LOFAR data[178].

8.4 Future Work

This work focuses on the substantial gains possible by parallelization of LOFAR data processing. We take in mind the complexity of our processing workflows, the full range of scientific pipelines and the heterogeneous nature of the underlying infrastructure. Because of these factors, a wide range of astronomical pipelines can use the software presented in this work. Future automation of the LoTSS processing requires deciding on data quality requirements at each step and automated re-processing strategies in case a data quality check fails. Implementing intelligent re-processing strategies will reduce the human supervision currently necessary to provide high-quality large-scale surveys such as LoTSS.

Scientific projects with significant data rates such as Gaia and LSST provide users with an integrated environment to efficiently process archived observations. Having such an environment is a necessary step to gain fast and easy to gain insights into LOFAR data. This work presents a method enabling scientists to incorporate processing hosted at scientific institutions and cloud providers to scale scientific processing horizontally.

One application for such large scale distributed processing is the Square kilometre Array. The Square Kilometer Array, (SKA) is a planned aperture synthesis radio telescope expected to have a total collecting area of one square kilometre. It is expected to produce more than 160 TB per day [179], data which needs to be processed. Scaling our tools to SKA-size processing requires a federation of clusters able to handle a high throughput workload. Nevertheless, as the SKA data processing will use different software tools than LOFAR, further study is needed on the optimal processing strategy for each of the many SKA science projects.

In recent years, academia has begun focusing on ease of access and reproducibility of science. Science done with cutting edge instruments, such as LOFAR, tends to be time-consuming to reproduce. Barriers such as setting up a working software environment and downloading massive data-sets prevent scientists from quickly and easily reproducing the results of their peers. These barriers make it difficult to verify the accuracy of discoveries and need to be overcome in order to make astronomy more honest and transparent. Automatically building, testing, versioning and releasing docker and singularity images is critical to making science in radio astronomy easily reproducible. Creating a LOFAR science portal and integrating our tools and software images with this portal is crucial to making radio astronomy both more accessible and more authoritative.

