



Universiteit
Leiden
The Netherlands

Characterising Radio Telescope Software With the Workload Characterisation Framework

Grange, Y.G.; Lakhoo, R.; Petschow, M.; Wu, C.; Veenboer, B.; Emsley, I.; ... ; Pasian, F.

Citation

Grange, Y. G., Lakhoo, R., Petschow, M., Wu, C., Veenboer, B., Emsley, I., ... Mariani, G. (2019). Characterising Radio Telescope Software With the Workload Characterisation Framework. *Astronomical Data Analysis Software And Systems Xxvi Asp Conference Series, Vol. 521, Proceedings Of A Conference Held (16-20 October 2016) At Stazione Marittima Conference Centre, Trieste, Italy. Edited By Marco Molinaro, Keith Shortridge, And Fabio Pasian. San Francisco: Astronomical Society Of The Pacific, 2019, P.683, 683*. Retrieved from <https://hdl.handle.net/1887/85249>

Version: Accepted Manuscript

License: [Leiden University Non-exclusive license](#)

Downloaded from: <https://hdl.handle.net/1887/85249>

Note: To cite this publication please use the final published version (if applicable).

Characterising radio telescope software with the Workload Characterisation Framework

Y. G. Grange¹, R. Lakhoo^{2*}, M. Petschow¹, C. Wu³, B. Veenboer¹, I. Emsley²,
T. J. Dijkema¹, A. P. Mechev⁴, and G. Mariani⁵

¹ *ASTRON, The Netherlands; (grange@astron.nl)*

² *Oxford e-Research Centre, University of Oxford, United Kingdom*

³ *ICRAR, University of Western Australia, Australia*

⁴ *Leiden Observatory and LIACS, Leiden University, The Netherlands*

⁵ *IBM Research, The Netherlands*

Abstract. We present a modular framework, the *Workload Characterisation Framework* (WCF), that is developed to reproducibly obtain, store and compare key characteristics of radio astronomy processing software. As a demonstration, we discuss the experiences using the framework to characterise a LOFAR calibration and imaging pipeline.

1. Introduction

Modern low-frequency radio interferometers, such as the LOw Frequency ARray (LOFAR) (van Haarlem et al. 2013), consist of many antennas and data processing is performed by software; this is reflected by the community term *software telescope*. Such software telescopes are modular and under constant development by large interdisciplinary research collaborations, working across continents; consequently, obtaining a holistic view of the entire system becomes a challenging task. It is a considerable challenge to manage, process and store such large datasets, within budgets and constraints, while “pushing the envelope” of technology. To achieve such cost-performance optimum, the first aim is to obtain a quantitative understanding of the compute, energy and data access behaviours exhibited by various radio astronomy data processing software pipelines and algorithms. Using low-level Linux kernel and hardware interfaces, the *Workload Characterisation Framework* (WCF) can be used to support operations and software development.

To encourage best practices, the work undertaken requires to be reproducible: the software should be open and publicly available, and it should allow for others to reproduce the software environment. In this work, we present the WCF that aims to provide key metrics to characterise workloads in a standard and reproducible manner. The WCF is designed to be extensible and we present an example of such an extension that aims to provide further insight when identifying software bottlenecks for research and development (R&D) purposes.

2. The Workload Characterisation framework

The WCF is under development as part of the Local Monitoring and Control (LMC) work package of the SKA Science Data Processor (SDP) consortium. The first purpose of the WCF is to provide the SDP compute resource scheduler with the essential workload characteristics for each pipeline processing component, in order to make optimal (cost-effective) scheduling decisions. Secondly, the WCF aims to assist telescope developers and operators to obtain a quantitative understanding of the compute, energy and data access behaviours exhibited by various pipelines, components and algorithms. Finally, the WCF aims to enable the micro-benchmarking of different compute platforms, which can be used for optimisation and comparison.

The WCF prototypes have already been used, or tested for SKA precursor and pathfinders, including LOFAR, MWA, JVLA HDR, ASKAPSoft and the CHILIES project.

3. The LOFAR use case

The WCF can be used for a variety of tasks; we focus on a bottleneck analysis – with the aim to support more targeted system development efforts (software and hardware). In addition to the typical WCF output data to assist this task, we consider non-time-series data to create the link between system behaviour and the structure of the software.

In this section, as an example use case, we consider a calibration and imaging processing pipeline, hereafter referred to as CALIB, which creates sky images from the LOFAR telescope.

3.1. The CALIB pipeline

The CALIB pipeline has been used to image HBA commissioning data of the Galactic diffuse synchrotron emission (Iacobelli et al. 2013).

For each frequency band, there are two pointings: one towards the target field and the other towards the calibrator field. As depicted in Fig. 1, the two input measurement sets are fed into the pipeline. After a copy operation, the calibrator is used to solve for the antenna gain amplitudes, which are then applied to the target. A phase calibration is then performed on both the calibrator and the target data. Finally, sky images are derived from the calibrated data sets. The entire calibration is performed using LOFAR’s DPPP (data preprocessing pipeline); the imaging is performed using WSClean (Offringa et al. 2014).




Figure 1. Work flow of the CALIB pipeline.

3.2. Goal-oriented system development

With goal-oriented or targeted system development, we promote the following work flow: (1) for a given micro-benchmark and platform, evaluate the baseline performance of both the hardware and software – ideally by replicating the results of measurements stored in a data base; (2) use the WCF to analyse system behaviour and identify bottlenecks; (3) concentrating on a specific part and possibly using other performance analysis tools, optimize the code or change the underlying hardware; (4) rerun the WCF measurements of the entire pipeline to assess the impact of the changes and store the results in a database.

With the WCF we assist a goal-oriented development in the following way: (1) Evaluating a pipeline’s performance over a long time requires a standardised storage format of the measurements. (2) By standardising the measurement tools, the measurements become comparable over a wide range of use cases: comparing different algorithms and their implementations, different hardware, and different input data and data formats. Among other advantages, this allows for meta-analysis of stored measurements for various configurations. (3) Quickly identify computational and other resource bottlenecks. This information could be used to alleviate software bottlenecks and identify the most salient system features that determine performance. Besides facilitating such a hardware-software co-design procedure, the insights can be used to inspire general research.

4. Results




Figure 2. Time series of system characteristics.

In Fig. 2 we present time series data of a CALIB pipeline run. We choose the metrics CPU and memory usage, memory bandwidth, and disk I/O bandwidth – all important parameters to identify performance bottlenecks. While a detailed analysis of the data is beyond the scope of this paper, we demonstrate its use for one specific step of

the pipeline, PHASESOLVE. This step requires roughly 55% of the entire execution time (Fig. 3; left). Further analysis indicates that the step exhibits poor scaling behaviour: the execution using a 28 core dual-socket Intel Xeon compute node is only 1.23 times faster than a sequential execution (Fig. 3; right). Using the OeRC SKA testbed, low-level CPU characteristics were gathered, and showed less than optimal CPU core usage on average, a high number of CPU migrations (292/sec) and context switches (3,521/sec). The poor scaling characteristics were verified by configuring OMP affinity settings, using a Round-robin (RR) real-time CPU scheduler and CPU pinning, only resulted in a 8.2% increase in runtime using 10 fewer CPU cores than the default. A natural next step is to improve on the scalability of the PHASESOLVE step of the pipeline.




Figure 3. Breakdown of execution time and speedup due to parallel execution.

5. Conclusion

Using a LOFAR calibration and imaging pipeline, we demonstrated the use of the WCF for R&D radio telescope software and hardware development. Such structured approach is especially relevant for the development of new instruments and algorithms for the SKA.

Acknowledgments. We thank Marco Iacobelli for providing the CALIB processing pipeline and the input data. This work is supported by the Dutch Ministry of EZ and the province of Drenthe through the ASTRON-IBM Dome grant and the EU FP7 under grant no ICT-610476 (DEEP-ER). The authors acknowledge the Department of Physics and the OeRC at the University of Oxford, for use of their system resources and personnel. The authors would like to acknowledge the use of the University of Oxford Advanced Research Computing facility in carrying out this work.

References

- Iacobelli, M., et al. 2013, A&A, 558, A72. 1308.2804
- Offringa, A. R., et al. 2014, MNRAS, 444, 606. 1407.1943
- van Haarlem, M. P., et al. 2013, A&A, 556, A2. 1305.3550