Vienna Correlator Report

Jakob Gruber[1], Frédéric Jaron[2], Johannes Böhm[2]

Abstract The Vienna Correlator is operated by the research unit Higher Geodesy of the Technische Universität Wien (TU Wien) and the Federal Office of Metrology and Surveying (BEV). The Vienna Scientific Cluster (VSC), a supercomputer located in Vienna, is used as a hardware component. A bandwidth of 10 GBit/s (shared by all VSC users), 480 compute cores, and 1 PByte memory for VLBI correlation are available. Two people are actively involved in VLBI correlation: J. Gruber and F. Jaron. In 2023, five out of a total of 35 (14%) VGOS-OPS sessions and, in 2024, five out of a total of 31 (16%) were correlated in Vienna. Besides processing VGOS sessions operationally, we develop new software solutions for VLBI data processing and simulation and use the Vienna Correlator for research activities. Our research and development activities are:

- Development and maintenance of the Vienna Raw Data Simulator for VLBI (VieRDS),
- Development of a wrapper software for integrating all processes in VLBI raw data processing, VLBI Integrator (VIN),
- Investigation of the effect of extended source structure in VGOS fringe-fitting,
- Investigation of signal properties for high-precision VLBI satellite missions, e.g., Genesis, using VieRDS.

1. Federal Office of Metrology and Surveying
2. Technische Universität Wien

VIEN Correlator

IVS 2023+2024 Biennial Report

1 General Information

The Vienna Correlator is operated by the research unit Higher Geodesy of the Technische Universität Wien (TU Wien) and the Federal Office of Metrology and Surveying (BEV). It became an official IVS correlator center in 2018. We correlate IVS VLBI sessions on an operational basis and for specific scientific projects. In particular, we contribute to the IVS by correlating the next-generation VLBI Global Observing System (VGOS) sessions. As a principal hardware component, we use a computer cluster called Vienna Scientific Cluster-4 [1] (VSC-4, see Figure 1). It is located next to our TU Wien offices (2.5 km away) and can be accessed remotely.

2 Component Description

For the correlation of VLBI level-0 data, we use the supercomputer resources of the VSC. The high-performance cluster called VSC-4 was installed in the summer of 2019. In total, the VSC-4 consists of 790 nodes. At the research unit Higher Geodesy, we have exclusive access to 10 of these nodes, which are reserved for projects related to VLBI correlation and are accessible anytime for the members of our VLBI group. In addition, storage of 1 PByte has been purchased to complete the VLBI correlation's hardware system, given the enormous data volumes in the VGOS era.

[1] https://vsc.ac.at/systems/vsc-4/

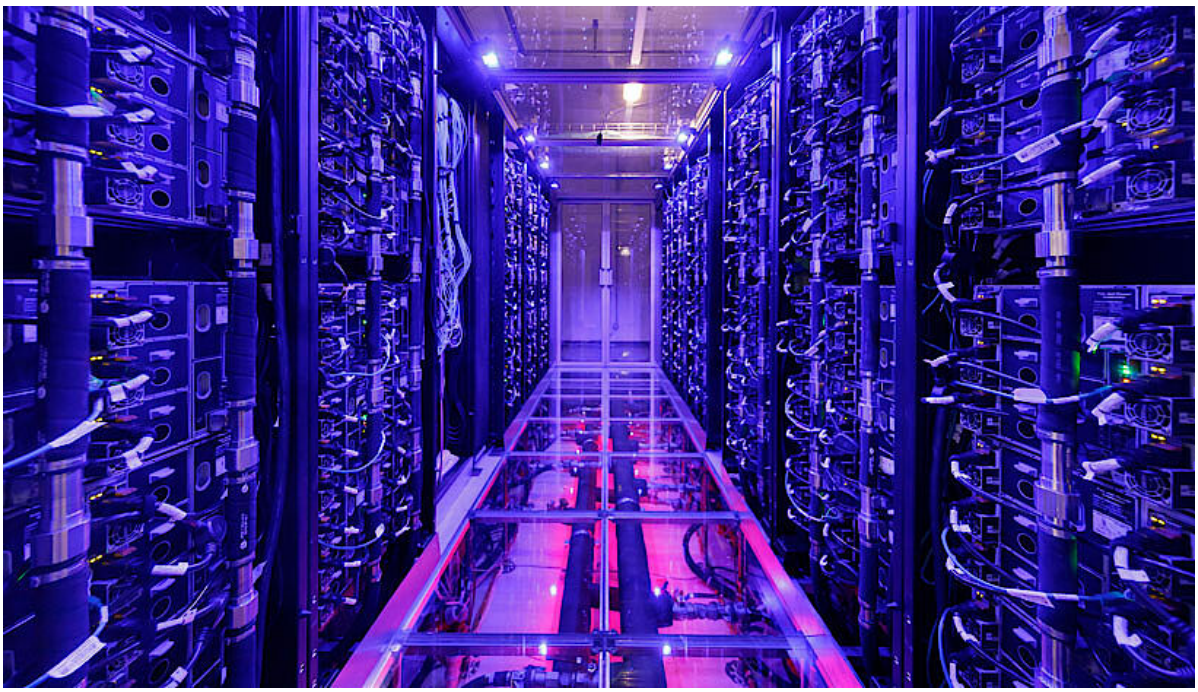

Fig. 1 VSC-4. 10 nodes and 1 PByte are reserved for VLBI correlation. Linked with 10 Gbps to the global research network GEANT. (©http://derknopfdruecker.com/)

2.1 VSC-4 Compute Cores and Storage

The VSC-4 is equipped with water-cooled Lenovo SD650 nodes, each with two Intel Skylake Platinum 8173 processors with 24 cores, interconnected with 100 Gbit/s OmniPath. Each high-performance node reaches 2.7 PFlops/s and has a main memory of 96 GBytes. For more information on the technical specifications of the VSC-4, please see https://vsc.ac.at//systems/vsc-4/.

By having access to 10 of these nodes, we can utilize 480 cores for VLBI correlation. For data storage, a General Parallel File-System (GPFS) with 1 PByte size is mounted to the VSC-4. This data volume is dedicated to VLBI correlation only within the VSC-4.

Currently, up to 15 VLBI stations are part of the VGOS observing program, resulting in approximately 420 TBytes for a 24-hour session. The Vienna VLBI Correlator has the hardware resources, in terms of the number of cores and storage size, to allow for correlation of the currently used and increasing VGOS station network in the future.

2.2 Software Capabilities

For the electronic data transmission between VLBI stations and the Vienna Correlator, we use the VLBI-specific e-transfer software jive5ab [2] and etd/etc [3] developed by H. Verkouter. Additionally, the high-speed network e-transfer program tsunami [4] is also installed on the VSC-4 login nodes and applied for operational level-0 data transmission. The Distributed FX style correlation software (DiFX, [1]) is installed to realize level-0 processing of the raw VLBI telescope data. While we continuously keep our DiFX installation up to date with the latest official releases, we also keep several older versions of DiFX to allow processing of VGOS and legacy S/X observations. The Slurm[5] workload management and job scheduling software is used to efficiently process the VLBI raw data by DiFX and reach a high parallelization level. Besides DiFX, the Haystack Postprocessing System (HOPS), PIMA, Pol-

[2] https://github.com/jive-vlbi/jive5ab

[3] https://github.com/jive-vlbi/etransfer

[4] https://sourceforge.net/projects/e-vlbi/files/tsunami.py

[5] https://slurm.schedmd.com/documentation.html

Convert, and nuSolve are installed at the VSC-4 to complete the entire raw data VLBI processing chain. Consequently, it is possible to process raw VLBI data and provide vgosDb files to the IVS community.

2.3 e-Transfer Performance

The Vienna Correlator is e-transfer only. Hence, the reception of hard disks containing VLBI level-0 data is not supported. The VSC-4 consists of ten login nodes linked to the GEANT network, allowing a maximum data rate of 10 Gbit/s. The 10 Gbit/s bandwidth is shared between login nodes and other users of the VSC. Several performance tests show a limit of 1.5 Gbit/s for a single e-transfer stream for most of the stations. The reason is unknown at the time of writing but represents an important topic for further investigation. However, multiple streams can be received by the Vienna Correlator in parallel, allowing for accumulated data rates up to 10 Gbit/s.

2.4 Correlation Performance

Slurm allows the parallelization of a single session, and several sessions can be processed in parallel. An investigation of the data throughput by DiFX on the VSC-4 showed excellent scaling with increasing the number of processing cores. For the correlation of a VGOS 24-hour session including up to nine telescopes, an optimal number of five nodes could be identified [4]. By using more than five nodes, the DiFX performance becomes data-limited rather than CPU-limited. Monitoring the total DiFX processing time of VGOS 24-hour sessions showed values between 21 and 24-hours [4]. The variation occurs due to varying level-0 data amounts. The maximum data throughput achieved by using 480 cores was 320-Gbps.

3 Staff

Two persons are involved in the work at the Vienna Correlator. Their names and most important responsibilities are listed below. Additionally, Johannes Böhm is a responsible contact point for the VSC-4 team.

- Jakob Gruber
  - Researcher
  - Maintenance of data transfer
  - VGOS correlation
  - Raw data simulation
  - Development of a wrapper software for integrating all processes in VLBI raw data processing (VIN)
- Frédéric Jaron
  - Postdoc researcher
  - Maintenance of data transfer
  - VGOS correlation
  - EU-VGOS correlation and organization
  - Special interest in source structure for VGOS
  - Investigation of novel methods for VGOS calibration
  - Development of third-party software to support correlation and fringe-fitting and correlation of various other special VLBI sessions

4 Current Status and Activities

4.1 VGOS Correlation

The Vienna VLBI Correlator is one of the permanent components of the International VLBI Service for Geodesy and Astrometry (IVS). The primary goal of the Vienna VLBI Correlator is to provide excellent VLBI correlation products with a fast turnaround time (TAT) in the VGOS era in accordance with the IVS quality standards.

The tasks for operating the Vienna VLBI Correlator include:

- Installation, maintenance, updating, and testing of all software packages required for VLBI correlation on the Vienna Scientific Cluster.
- Implementation, contact management, and testing of electronic data transfer between the Vienna VLBI Correlator and the VLBI station network used for observations.

- High-performance processing of raw data transmitted by VLBI stations on the Vienna Scientific Cluster to ensure a fast TAT.
- Exception handling, data flagging, and quality assurance of the obtained VLBI correlation products in compliance with IVS quality standards.
- Uploading VLBI correlation products to the VLBI data server of the National Aeronautics and Space Administration (NASA).
- Participation in international VLBI meetings to discuss current VLBI correlation-related challenges, develop solutions, and respond promptly to decisions made by the IVS Directing Board.

4.2 Development of a Wrapper Software for Integrating All Processes in VLBI Raw Data Processing – VLBI Integrator (VIN)

The VLBI Global Observing System (VGOS) aims for continuous 24/7 VLBI observations with fast turnaround times for geodetic products. However, a major bottleneck in achieving these goals is the correlation process, including data transfer.

An in-house project focuses on automating the correlation process and data transfer by integrating all necessary and existing software packages. The resulting software, called VLBI Integrator (VIN), serves as a wrapper software to streamline all processes related to VLBI raw data processing. The aim is to process raw VLBI data more automatically and work towards the bold VGOS goals.

4.3 Source Structure in VGOS Fringe-Fitting

At TU Wien we are working on a research project concerning source structure mitigation in level-1 geodetic VLBI data. The objective of this project is to correct the output visibilities of the DiFX software correlator so that in the ideal case the data look like they had been obtained by the observation of point sources. We use the Vienna Scientific Cluster to apply our method to observational data from a selection of VGOS sessions. We then post-process the data in the conventional way and obtain vgosDb files, which we then analyze with regard to closure quantities, also performing a full geodetic analysis. In this way, we can quantify the impact that the source structure and its mitigation have on observables and their analysis. Preliminary results from this project have been documented in the recent proceedings of the EVN Symposium [2].

4.4 The Genesis Science Exploitation Team (GSET) – Working Group 3 (WG-3)

The Genesis Science Exploitation Team (GSET) of the European Space Agency (ESA) consists of several working groups (WG). Working Group 3 (WG-3) focuses on the VLBI system and has the following key responsibilities:

- Advising and supporting ESA in the design of the Genesis VLBI transmitter.
- Investigating and ensuring compatibility with regular IVS operations and IVS product generation.

We are actively contributing to this working group by developing software solutions and conducting raw data simulations with VieRDS. Using established data processing software packages installed on the Vienna Correlator such as DiFX and HOPS, we test new Genesis observing modes. Our main focus is on analyzing potential signal properties emitted by Genesis to ensure high-accuracy geodetic VLBI.

5 Future Plans

Moving forward, we will continue to focus sharply on VGOS correlation, ensuring our active participation and correlation of VGOS-OPS sessions. Based on our experience from the past years in correlating VGOS sessions, and with the upcoming VGOS sessions in 2025 and 2026, we will further develop VLBI Integrator (VIN) to improve turnaround time (TAT) and consistency of VGOS session processing.

Additionally, we will leverage the capabilities of the Vienna Correlator to simulate raw data for the Gen-

esis mission, enhancing our ability to test and validate new observation modes.

By 2025, we will also begin using the new realization (5) of the VSC system, ensuring we maintain and expand our computing power to meet the increasing challenges in VLBI raw data processing over the coming years.

Acknowledgements

The authors are grateful to the Austrian Science Fund (FWF) for supporting this work with the project P31625 (VGOS Squared). The computational results presented have been achieved in part using the Vienna Scientific Cluster (VSC).

References

1. A. T. Deller, Brisken W. F., Phillips C. J., Morgan J., Alef W., Cappallo R., Middelberg E., et al., 2011, DiFX-2: A More Flexible, Efficient, Robust, and Powerful Software Correlator, PASP, 123, 275. doi:10.1086/658907.
2. F. Jaron, Baldreich L., Böhm J., Charlot P., Collioud A., Gruber J., Krásná H., et al., Mitigating Source Structure in Geodetic VLBI on the Visibility Level, Proceedings of the 16th EVN Symposium, Ed. E. Ros, P. Benke, S. A. Dzib, I. Rottmann, & J. A. Zensus, Bonn: Max-Planck-Institut für Radioastronomie, 2024, pages 145–148.
3. J. Gruber, J. Böhm, A. Nothnagel (2021) VieRDS: A Software to Simulate Raw Telescope Data for Very Long Baseline Interferometry, *Publications of the Astronomical Society of the Pacific*, doi:10.1088/1538-3873/abeca4.
4. J. Gruber (2022) The Vienna VLBI Raw Data Simulator and Correlator in the VGOS Era, *Dissertation, Technische Universität Wien*.