Vienna Correlator Report 2021–2022

Jakob Gruber[1], Frédéric Jaron[1], Leo Baldreich[1], Johannes Böhm[1]

Abstract The Vienna Correlator is operated by the research unit Higher Geodesy of the Technische Universität Wien (TU Wien). 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. Three people are actively involved in VLBI correlation: J. Gruber, F. Jaron, and L. Baldreich. In 2021, seven out of a total of 27 (26%) VGOS-OPS sessions, and in 2022, 12 out of a total of 44 (27%) were correlated in Vienna. In the EU-VGOS project, sessions were observed, correlated, analyzed, and new software solutions were organized and developed.

1 General Information

The Vienna Correlator is operated by the research unit Higher Geodesy of the Technische Universität Wien (TU Wien). 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 the Vienna Scientific Cluster-4[1] (VSC-4, see Figure 1). It is located 2.5 km away from our TU Wien offices and can be accessed remotely.

1. Technische Universität Wien, Department für Geodäsie und Geoinformation, Wiedner Hauptstraße 8, 1040 Wien, Austria

VIEN Correlator

IVS 2021+2022 Biennial Report

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

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 ten 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, 1 PByte of storage was purchased to complete the VLBI correlation's hardware system, given the enormous data volumes in the VGOS era.

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 8,173 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 GByte. For more information on the technical specifications of the VSC-4, please see https://vsc.ac.at//systems/vsc-4/.

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

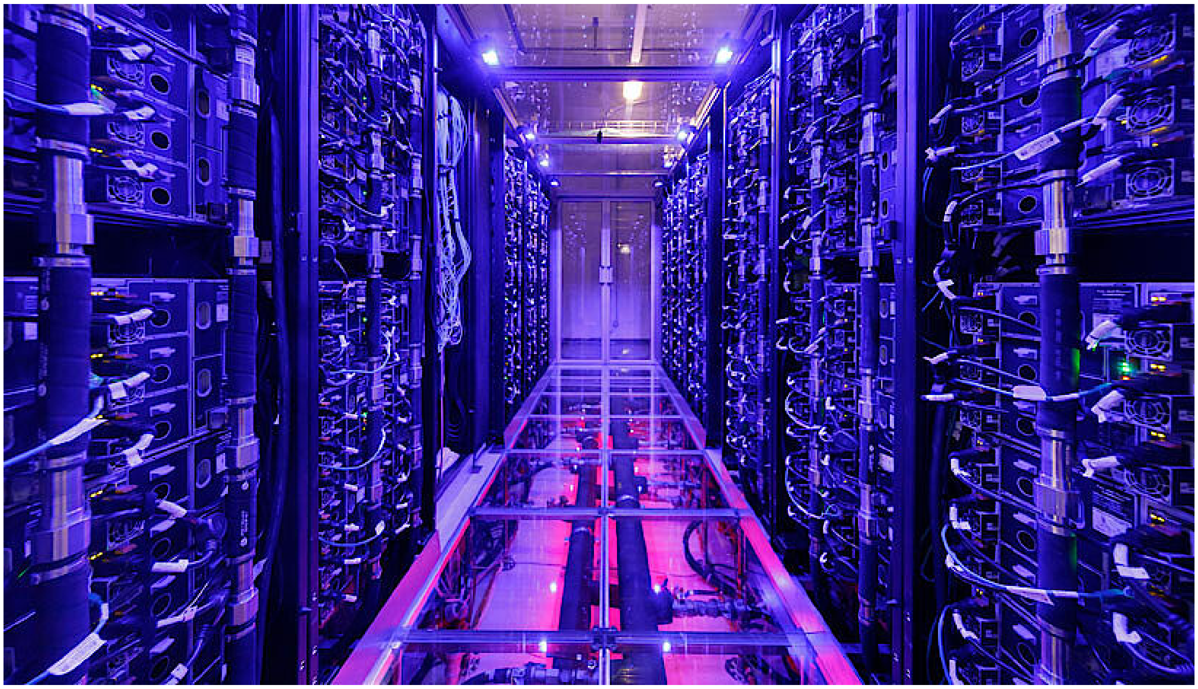

Fig. 1 VSC-4. Ten nodes and 1 PByte are reserved for VLBI correlation. Linked with 10 Gbps to the global research network GEANT. (©http://derknopfdruecker.com/)

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, [3]) 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 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, PolConvert, 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] `https://github.com/jive-vlbi/jive5ab`

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

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

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

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 [10]. The reason is unknown at the time of publication but represents an important topic for further investigation.

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 [10]. 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 [10]. The variation occurs due to varying Level-0 data amounts. The maximum data throughput achieved by using 480 cores was 320 Gbps.

3 Staff

Three people are involved in the work at the TU Wien IVS VLBI correlator. Their names and most important responsibilities are listed below. Additionally, Johannes Böhm is a responsible contact point for the VSC-4 team, and Axel Nothnagel acts as a consultant concerning correlation/fringe-fitting and raw data simulation.

- Jakob Gruber
  - Postdoc researcher
  - Maintenance of data transfer
  - VGOS correlation
  - Raw data simulation
  - Development of third-party software to support correlation and fringe-fitting and correlation of various other special VLBI sessions
- Frédéric Jaron
  - Postdoc researcher
  - Maintenance of data transfer
  - VGOS correlation
  - EU-VGOS correlation and organization
  - Special interest in source structure for VGOS
  - Development of third-party software to support correlation and fringe-fitting and correlation of various other special VLBI sessions
- Leo Baldreich
  - Student Assistant
  - Maintenance of data transfer

4 Current Status and Activities

4.1 VGOS Correlation

The next-generation VGOS system became operational, beginning regular observing 24-hour sessions in 2020. The provision of high-quality VGOS raw data products by the Vienna VLBI correlator was examined in two comparison studies conducted by the VLBI group of the MIT Haystack Observatory in 2020. The verification showed that the Vienna VLBI correlator could provide high-quality products. As a result, the Vienna VLBI correlator became an operational IVS component for VGOS correlation in 2021. In total, there were 27 VGOS 24-hour sessions conducted this year. Seven were correlated by the Vienna VLBI correlator, eight by the Bonn correlator, six by the Washington correlator, four by the Shanghai correlator, and two by the MIT Haystack correlator. In this regard, the Vienna VLBI correlator contributed to a successful launch of operational VGOS with the correlation of 25% of all VGOS sessions in 2021 and 27% in 2022, respectively, making it an essential part of the global VGOS infrastructure.

4.2 EU-VGOS

The EU-VGOS project was initiated in 2019 with the aim of investigating methods for correlation and post-processing of VGOS data [1]. A recent description of the project's current status can be found in [2]. At the time of writing, EU-VGOS is a collaboration of about 50 individuals working at different institutes, mainly in Europe but also beyond. The project is structured into five Working Groups (WGs): 1) Management, 2) Operations, 3) e-Transfer, 4) Correlation, and 5) Analysis. Three of these WGs are led by employees of TU Wien: WGs 1 and 4 by Frédéric Jaron, and WG 3 by Jakob Gruber.

4.2.1 WG e-Transfer

The e-transfer of observational data from the stations to the correlator is still one of the most critical bottlenecks of the VGOS processing chain. It still takes significant time for the data to arrive at the correlator. Furthermore, data rates are often much below the expected values. For these reasons, the EU-VGOS collaboration established a dedicated WG for this topic. One quickly identified problem was that relatively small files (i.e., several GByte each) produce a significant overhead for e-transfer tools that significantly reduce the effective data rate. The etd/etc software overcomes this problem and can speed up the transfer. This is especially relevant when transferring multifile data. This was demonstrated in [4], where the transfer duration could be reduced from 18 down to 12 hours in one example.

4.2.2 WG Correlation

At the beginning of 2021, the correlation of VGOS data with the DiFX software correlator required several patches to be applied to the software. This quickly led to inconsistencies in the VGOS correlation software between the different Correlator Centers. This issue was discussed in the EU-VGOS Correlation meetings. As a result of these discussions, the DiFX version 2.5.4 was developed, tested by the larger DiFX user community, and eventually officially released.

The WG Correlation investigates the application of PolConvert [5, 6] to VGOS data. The main functionality of PolConvert is to convert the correlator output of linearly polarized data to circular polarization. But in order to do so, the software first estimates the complex cross-polarization gain differences between the two linear feeds of each antenna in terms of amplitude and phase. PolConvert thus offers a possibility to combine the data to Stokes I for fringe-fitting.

The session er2201 was observed in September 2022. The background of this session was that the source 4C 39.25 was identified as an ideal cross-polarization gain calibrator during the analysis of previous EU-VGOS sessions. The session was correlated at TU Wien (and also at MIT Haystack Observatory). At TU Wien, we are analyzing the data. Results so far have shown that the cross-bandpasses of the participating stations could be estimated with excellent quality, and it is possible to investigate the time evolution of the cross-bandpasses over the six-hour duration of the session. We can already report that the principal shape of the bandpasses remains stable, but there is also perceivable variability. More details will be available in an upcoming publication.

4.3 Raw Data Simulation

As part of J. Gruber's dissertation, a novel tool for simulating VLBI raw data was developed [9]. The Vienna Raw Data Simulator (VieRDS) can be freely downloaded from Github[6] and is available to the whole IVS community. VieRDS was successfully used to study the quality of the phase calibration signal as a function of injection power. In addition, we investigated how the characteristic station frequency response affects the signal-to-noise ratio of the observed quantity. In cooperation with H. Hase, the quality of the group delay was investigated using alternative frequency sequences [8]. In summary, VieRDS has proven to be very useful for studies of effects that cannot be simulated at the group delay level.

5 Future Plans

In the near future, we will continue to focus sharply on VGOS correlation and the organization of the EU-VGOS project. Specifically, we will correlate six VGOS-OPS sessions in 2023, with the optimization of correlation for more efficient processing of raw VLBI data as an important goal. In addition, we plan to work intensively on source structure effects for improved VGOS. The tool VieRDS developed by us will be used for this purpose.

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 were achieved in part using the Vienna Scientific Cluster (VSC).

[6] `https://github.com/TUW-VieVS/VieRDS`

References

1. Alef W., Anderson J. M., Bernhart S., de Vicente P., González García J., Haas R., La Porta L., et al. (2019) The EUropean-VGOS Project, *Proceedings of the 24th European VLBI Group for Geodesy and Astrometry Working Meeting, 17-19 March 2019, Las Palmas de Gran Canaria, Spain*, Eds. R. Haas, S. Garcia-Espada, and J. A. López Fernández, ISBN: 978-84-416-5634-5, pp. 107–111.
2. Albentosa E., Alef W., Bernhart S., Böhm J., Bolaño González R., Choi Y. K., Gansmoe T., et al. (2023) Current Status of the EU-VGOS Project, *International VLBI Service for Geodesy and Astrometry 2022 General Meeting Proceedings*, Eds. Kyla L. Armstrong, Dirk Behrend, Karen D. Baver, NASA/CP-20220018789, pp. 85–89.
3. Deller A. T., 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
4. Jaron F., Bernhart S., Böhm J., González García J., Gruber J., Choi Y. K., Martí-Vidal I., et al. (2021) EU-VGOS activities in Vienna, *Proceedings of the 25th European VLBI Group for Geodesy and Astrometry Working Meeting, 14-18 March 2021 Cyberspace, Gothenburg, Sweden*, Ed. R. Haas, ISBN: 978-91-88041-41-8, pp. 19–23.
5. Martí-Vidal I., Roy A., Conway J., Zensus A. J. (2016) Calibration of mixed-polarization interferometric observations. Tools for the reduction of interferometric data from elements with linear and circular polarization receivers, A&A, 587, A143. doi:10.1051/0004-6361/201526063
6. Martí-Vidal I., Mus A., Janssen M., de Vicente P., González J. (2021) Polarization calibration techniques for the new-generation VLBI, A&A, 646, A52. doi:10.1051/0004-6361/202039527
7. J. Gruber, F. Jaron, A. Nothnagel, J. Böhm (2021) Vienna Correlation Center Biennial Report 2019/2020, *International VLBI Service for Geodesy and Astrometry 2019+2020 Biennial Report*, Edited by Dirk Behrend, Kyla L. Armstrong, and Karen D. Baver, NASA/TP–20210021389.
8. J. Gruber, J. Böhm, F. Jaron, A. Nothnagel (2021) Simulation of VGOS observations at the raw data level with VieRDS, *Proceedings of the 25th European VLBI Group for Geodesy and Astrometry Working Meeting, 14-18 March 2021 Cyberspace, Gothenburg, Sweden*, Ed. R. Haas, ISBN: 978-91-88041-41-8, pp. 28–33.
9. 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.
10. J. Gruber (2022) The Vienna VLBI Raw Data Simulator and Correlator in the VGOS Era, *Dissertation, Technische Universität Wien*.