Status Report of the Bonn Correlator

Simone Bernhart[1,2,3], Yoon Kyung Choi[1,2,3], Helge Rottmann[3], Jan Wagner[3], Christian Plötz[2]

Abstract We present a status report for the Bonn Correlation Center for 2022 and 2023 focusing on geodesy. We summarize our duties as one of the IVS correlators and the experience we have gained in the correlation of VGOS observations.

Keywords VLBI correlator, DiFX

1 Introduction

The Bonn correlator, located in Bonn, Germany, is operated jointly by the Max Planck Institute for Radio Astronomy (MPIfR) in Bonn and the Federal Agency for Cartography and Geodesy (Bundesamt für Kartographie und Geodäsie, BKG) in Frankfurt. The MPIfR hosts the correlator facility and provides the Internet connectivity. Since January 2017, the personnel responsible for the correlation of geodetic sessions are employed by the BKG via a private contractor, the Reichert GmbH. The costs of the cluster hardware are shared between MPIfR and BKG. Moreover, the MPIfR personnel support the geodetic staff in technical, hardware and software issues.

The Bonn correlator also serves as an inherent test bench for the DiFX and e-transfer software, so that all its personnel contribute to the debugging of these tools and some also to their further development.

1. Reichert GmbH
2. Federal Agency for Cartography and Geodesy (BKG)
3. Max Planck Institute for Radio Astronomy

2 Correlator Capabilities

The Distributed FX software correlator (DiFX, [1]) is used at the Bonn correlator in various versions. The correlator is running on a high-performance computing (HPC) cluster, which was renewed in 2015 to match both VGOS and mm-VLBI requirements. It consists of 68 nodes with 20 compute cores each, for a total of 1,360 cores; three head nodes which allow execution of several correlations and post-processing in parallel; 2.8 PB disk space in RAID units (Figure 1) as well as 14 Mark 5 playback units and 11 Mark 6 playback units of which five have four and six have six bays (Figure 2).

The raw data are recorded at the stations either on modules (Mark 5 or Mark 6) or on storage servers, usually referred to as Flexbuffs. For geodetic experiments the data are mostly e-transferred to the HPC cluster. The cluster has a 10-Gbps Internet connection.

Fig. 1 Correlator compute nodes and raids.

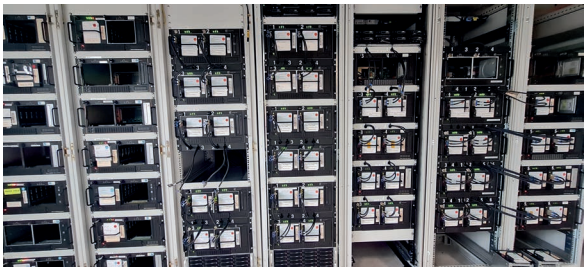

Fig. 2 Mark 5 and Mark 6 units.

3 Correlation Cluster

The RAID units are combined in a so-called BeeGFS, a hardware-independent POSIX parallel file system.

We use SLURM as a cluster workload manager where batch jobs are submitted for parallel processing, since the cluster is shared with other users in the MPIfR, e.g., for EHT or GMVA correlation, preprocessing for pulsar search or numerical simulations.

The correlator output data (SWIN files) can be exported to FITS and HOPS (Mark IV) formats. For postprocessing, the following software packages are available: AIPS, CASA, PIMA, and HOPS (Haystack Observatory Postprocessing System[1]), the latter of which is the standard tool for geodesy. The correlator outputs and other important files (e.g., VEX and v2d files) are backed up daily on the HPC cluster. The final products are archived on the MPIfR archive server, where they will be kept for at least ten years. The EXPAD and COMEDIA tools are used for bookkeeping of experiments and corresponding media correlated in Bonn. They are the frontends to a local database which records all relevant information such as the observation date, participating stations, modules, and status of the experiment.

4 e-Transfer to/from Bonn

Since October 2021, the cluster is connected to the Internet through a commercial 10-Gbps line. A firewall computer for the e-transfer servers was set up in late 2011. Moreover, a 1-Gbps connection still exists to the Bonn University as part of the German Research Network (Deutsches Forschungsnetz, DFN) in order to enable data transfers for those facilities that can only send data through national research networks (NRENs) such as Warkworth station. In January 2024, one of the aging e-transfer servers (eportal1) was replaced by new hardware.

Currently, the most commonly used transfer protocol for geodetic Internet transfers is jive5ab/m5copy[2]. Developed as a VLBI data recorder software, it allows fast and flexible data transfers through high-speed internet connections.

Since 2019, we have also been using the e-transfer server/client system (etc/etd)[3] where a client program runs server-to-server transfers. It supports the usage of remote wildcards superseding the login to a remote machine.

Fig. 3 More information on e-transfers with the Bonn correlator: https://www3.mpifr-bonn.mpg.de/div/vlbi/geodesy/evlbi/index.html.

The Bonn correlator still hosts the e-transfer web page (https://www3.mpifr-bonn.mpg.de/cgi-bin/showtransfers.cgi) that shows ongoing transfers between stations and correlators as well as the bandwidth and storage capacities that are still available. More information concerning data transfers with the Bonn correlator can be found via the QR code/URL in Figure 3.

[1] https://www.haystack.mit.edu/haystack-observatory-postprocessing-system-hops/

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

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

5 Software Versions

We are currently using DiFX-2.6.3 for the legacy S/X correlation and post-processing, which comes with HOPS v3.23. All legacy Rapid sessions have, in the meantime, turned into so-called mixed-mode sessions due to the participation of several (VGOS) stations, that record dual-linear polarization as opposed to the observing mode of the original legacy stations, which is right circular polarization (RCP). This requires the application of option "mixed_pol_yshift90 true" in the fourfit control file which only comes with the latest release of HOPS v3.25 and adds a 90 degree phase shift to VGOS–VGOS baselines to get a coherent pseudo-Stokes-I combination and avoid decorrelation of the signal.

VGOS correlation and post-processing is performed by means of DiFX-2.5.5 (HOPS v3.22) together with HOPS v3.24/v3.25 for VGOS-OPS sessions and HOPS v3.24 for VGOS-INT-Ms.

Before switching to a new DiFX version, the software output (fringe visibility measurements) needs to be compared thoroughly to the output from previous software versions. In order to do so, we use several tools. One is diffDiFX.py, which comes with each DiFX version and directly compares the correlator output, i.e., the SWIN files of each correlation run. Another tool is called compare-baselines-v6.pl; it compares the alists of the different fringe fitting runs. In a final step, we extract the observables by means of *aedit* and plot them against each other. The values should ideally lie on a 45° line.

6 Sessions Correlated in Bonn

Our duties include the correlation of Intensive series (IVS-INT-3), IVS-R1 series as well as IVS-OHIG and IVS-T2 series for IVS S/X legacy sessions and VGOS 24-hour sessions.

In 2022, we correlated 53 IVS-R1 (24 hours, weekly) sessions, eleven IVS-T2 (24 hours, bi-monthly) sessions, six IVS-OHIG (24 hours, bi-monthly) sessions, eleven VGOS-OPS (24 hours) sessions, and 31 IVS-INT-3 (one hour, weekly on Monday) sessions.

In 2023, we processed 51 IVS-R1 sessions, ten IVS-T2 sessions, eight IVS-OHIG sessions, eight VGOS-OPS sessions, and 21 IVS-INT-3 sessions.

Due to a failure of the SESHAN telescope in the last quarter of 2023, the weekly IVS-INT-3s have been suspended. As a replacement and future addition it was decided to observe the VGOS-INT-M series instead. Starting in October, we processed seven of these sessions in 2023.

7 Testing and Troubleshooting

At the Bonn correlator, we frequently perform test correlations of new recording systems at the stations. If, for example, an antenna wants to switch from Mark 5 to DBBC2/3 data recording, these two recording modes are used simultaneously and the subsequent correlation is performed in parallel. A test database is then submitted to the analysts for comparison who finally give their feedback to the station.

In April 2023, we started testing DiFX-2.8.1. However, we noticed that baselines of mixed polarization show higher amplitudes (factor of 4.7) and higher SNR than with DiFX-2.6.3, whereas baselines with the same polarization show the same numbers in both versions. Our fix for this difx2mark4 issue is pending verification by Haystack. For the time being, we have paused further testing.

An example of troubleshooting is the correlation of VGOS-OPS session vo4024. It was observed on 24 January 2024 by eleven stations in total. Three of the stations shipped their Mark 6 modules to the correlator, the remaining data was e-transfered. After the arrival of the last data on 2 February, the correlation was started on the same day. After 2.5 days, the correlation had seemingly finished. Unfortunately, it turned out that only 50% of the data was processed. The issue was caused by Open MPI (an open source Message Passing Interface implementation in DiFX), something that had already been reported by the AuScope colleagues. Indeed, Jan Wagner had recommended ordering the antennas in the v2d file such that Mark 6 modules/antennas are listed first, which was disregarded in the first correlation run. Moreover, he created a customized version of DiFX-2.5.5 with Open MPI vers. 5.0.0 (currently open MPI 4.1.6). It still took several passes with this adjusted DiFX version until the re-

maining scans were correlated successfully. Despite this problem and thanks to the speedy help of Jan, we managed to submit the database of this session within 21 days after observation.

8 Outlook

The correlation cluster hardware was made VGOS-ready in 2015, but has become susceptible to aging by now. Hence, a hardware upgrade is envisaged for the next year.

Moreover, the geodesy group will continue testing the latest DiFX version (currently DiFX-2.8) before applying it for normal operation. After comparing the SWIN files as well as the resulting observables of the presently used DiFX versions (DiFX-2.5.5 and 2.6.3) and the upcoming release, we will switch to the latter one as soon as possible—its stability presupposed.

References

1. A. T. Deller, W. F. Brisken, C. J. Phillips, J. Morgan, W. Alef, R. Cappallo, E. Middelberg, J. D. Romney, H. Rottmann, S. J. Tingay & R. Wayth, “DiFX-2: A More Flexible, Efficient, Robust, and Powerful Software Correlator”, PASP, 2011, 123, 275–287.