The Bonn Correlator: Status Report

Laura La Porta [1,2], Walter Alef [3], Simone Bernhart [1,2], A. Müskens [4], Axel Nothnagel [4], Helge Rottmann [3], Torben Schüler [2], Jan Wagner [3]

Abstract We present a status report of the Bonn correlator. After discussing some technical aspects concerning the cluster and its performance, we will introduce the people working at the correlator, as well as the ongoing projects and duties, focusing on geodesy.

Keywords VLBI correlation, DiFX, VGOS

1 Introduction

The Bonn correlator is operated jointly by the Max Planck Institute for Radioastronomy (MPIfR) in Bonn and by the Federal Agency for Cartography and Geodesy (Bundesamt für Kartographie und Geodäsie, BKG), with the support of the Institute of Geodesy and Geoinformation (IGG) of the Bonn University. The MPIfR hosts the correlator facility and shares with the BKG the costs of the cluster, of most of the staff and of the internet connectivity. The IGG contributes to the connectivity of the cluster and pays one member of the geodetic staff. Since January 2017 the personnel responsible for the correlation of geodetic sessions have been employed by the BKG via a private contractor, the Reichert GmbH.

1. Reichert GmbH
2. Bundesamt für Kartographie und Geodäsie
3. Max-Planck-Institut für Radioastronomie
4. Institut für Geodäsie und Geoinformation der Rheinische Friedrich-Wilhelms Universität Bonn

2 The HPC Cluster

We run the Distributed FX software correlator at Bonn ([Deller et al. 2011]), of which we have various versions at our disposal. In particular, a branch version developed by J. Anderson for Radioastron experiments exists ([Bruni et al. 2014]). For production we use the latest stable DiFX release, and before switching to a newer DiFX version we perform a comparison of the resulting observables. Currently we are using DiFX 2.5.2, for both geodesy and astronomy. The correlator is running on a High Performance Computing (HPC) cluster, which was renewed in 2015 to match both VGOS and mm-VLBI requirements and consists of:

- 68 nodes with 20 compute cores each, for a total of 1,360 cores, which provide a computing power about ten times larger than that available with the old cluster;
- three head nodes which allow execution of several correlations in parallel (up to three parallel correlations were tested, and no reduction in speed was observed);
- 56 Gbps Infiniband interconnect between all nodes;
- 1.348 PB of disk space organized in RAID units (each with redundancy), of which 1.2 PB are combined in a BeeGFS parallel cluster file system;
- 14 Mark 5 playback units; and
- eight Mark 6 playback units each with four bays.

The raw data are recorded at the stations on Mark 5 or Mark 6 modules or on flexbuff. For geodetic experiments the data are mostly e-transferred to the HPC cluster, connected to the Internet through two 1-Gbit lines, one of which belongs to the Bonn University. Various data formats have already been correlated in Bonn: Mk4, Mk5, DVP, and VDIF. A native playback

Fig. 1 New HPC cluster at MPIfR seen through a glass wall.

mode for Mark 6 modules is about 90% ready. A feature to correlate multiple datastreams per station was implemented and is now routinely used, e.g., for the Event Horizon Telescope (EHT) sessions.

The correlated data can be exported to FITS and HOPS (mk4) format. The post-processing may be done with AIPS, PIMA, and HOPS, 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 have been extended to help bookkeeping of the experiments correlated in Bonn by collecting all relevant information (observation date, participating stations, modules, status of the experiment) stored in the local database (difxdb).

3 Staff at the Bonn Correlator

The MPIfR staff at the Bonn correlator belongs to the VLBI technical development group, headed by W. Alef. Its members are H. Rottmann, A. Roy, J. Wagner, Y. Pidopryhora, M. Lisakov, S. Dornbusch, and G. Tuccari, who has a guest contract. In addition to the scientific staff, there are one technician, R. Märtens, one engineer, M. Wunderlich, and two operators, H. Sturm and H. Fuchs. The group is responsible for keeping the cluster software up to date, for hardware maintenance and repair, as well as for IT support and software correlator improvements (H. Rottmann and J. Wagner are DiFX developers).

The group members are involved in several projects. A. Roy is responsible for the Atacama Pathfinder EXperiment (APEX) and H. Rottmann for the beamformer of the Atacama Large Millimeter/submillimeter Array. G. Tuccari is the leader for the Digital Base Band Converters (DBBCs) and the Fila10G and, together with W. Alef, for the development of a new BRoad bAND (BRAND) receiver. J. Wagner, Y. Pidopryhora, and M. Lisakov take care of the correlation of astronomical experiments, which are focused on very high resolution imaging. About five Radioastron sessions are correlated in Bonn every year. Those sessions involve up to 38 antennas and baseline lengths of several Earth diameters. Two Global Millimeter VLBI Array (GMVA) sessions with up to 21 antennas are also correlated in Bonn. The data rate ranges between 2 Gbps and 8 Gbps, depending on

Fig. 2 View of the Mark 5 and Mark 6 units through a glass wall.

whether ALMA is participating in the observations or not, so that the amount of stored data can be as large as 700 TB. In addition, half a session of the Event Horizon Telescope (EHT) is correlated in Bonn, namely the data at 230 GHz (1 mm), which includes left and right circular polarization over a 4 GHz bandwidth. Last but not least, the group performs several tests for development of digital VLBI backends.

The geodesy group at the Bonn correlator is composed of about 2.3 FTEs. A. Müskens is mainly concerned with the schedules of various IVS sessions, namely of INT3, EURO, T2, and OHIG, which he generates with the SKED software ([Gipson 2010]). S. Bernhart and L. La Porta are responsible for the remaining activities. They coordinate the data logistics, prepare and supervise the correlation, carry out the post-processing and deliver the resulting observables to the IVS repository in the form of databases, which are suited for the subsequent geodetic analysis software packages. The IVS sessions correlated in 2017 are 45 R1, six EURO, six T2, 44 INT3, four OHIG, and five days of CONT17. Since May of this year those databases have been produced solely via the vgosDB-make software ([Bolotin et al. 2016]) in VGOS format.

Aside from these standard duties, they provide the stations with feedback on their performance and support tests of the VLBI systems, in particular for the Wettzell Observatory. They also participate in IVS activities by cooperating on various subjects, for example the determination of peculiar clock offsets for the various stations or the verification of the distributed correlation as a possible approach for sharing the work load among correlators during the VGOS era (read next section for further details).

As a final remark, the Bonn correlator is a natural test-bench for the DiFX software and for the e-transfer protocols, so that all its personnel contribute to debugging those tools.

4 Bonn as an IVS Correlator

4.1 CONT17

The Legacy-1 S/X network of the IVS CONTinous VLBI campaign 2017 was correlated in Bonn. Standard activities were stopped a couple of weeks before CONT17 to prepare storage space and organize the logistics of data transfers. We stored about 500 TB, which were mostly e-transferred to Bonn (only three stations sent modules). The cluster BeeGFS failed at the beginning of the campaign, so that we relied solely on three RAIDs for storing and correlating the Rapid-like sessions (C1701-R1, C1703-R4, C1707-R4, and C1714-R1). As a consequence the correlation ran much slower (by a factor of three) w.r.t. normal, and we had to pause e-transfers until the beginning of January when the cluster BeeGFS was completely restored. Nevertheless, we managed to submit the databases of the Rapid-like sessions within the usual two to three week latency time. The final correlation of CONT17 was completed during February.

It took some time before we resumed our normal activities (at the beginning of May), due to some doubts concerning the global set of clock parameters used for the final correlation. The following discussions concluded that the set of clocks was correct and no re-correlation was necessary.

The effective processing time of the CONT17 campaign was enormously reduced with respect to CONT14, thanks to the capabilities of the new cluster. The computing time for 24 hours of data was about a factor of three higher for CONT14, also due to the larger number of modules involved in the correlation (Mk5 units often had to be reset).

4.2 Distributed Correlation

A possible way to deal with the huge workload foreseen for VGOS could be to share it among several correlators by dividing the sessions into time blocks. Each correlator would receive only part of the raw data for a given session. Upon request of the IVS Directing Board we organized for testing such an approach, together with five other correlators (Onsala, Warkworth, Hobart, Seshan, and Vienna). We agreed on a common DiFX and HOPS version and performed the test for session R1840, which belongs to the IVS sessions regularly processed in Bonn.

As Bonn is the main correlator, we prepared and sent to our colleagues the vex and v2d files to be used for correlation, as well as the control file for fringe-fitting the data. We then collected the DiFX and Fourfit outputs of the branch correlators to compare them with ours. Verification of results is ongoing. We will generate a new database for R1840 by combining the output of the main and the branch correlators. R. Haas will perform a geodetic analysis of that database and compare the outcome with that of the original database, which contains only the outputs of the Bonn correlator.

The main downside of a distributed correlation is that the data logistics becomes more complicated. The raw data should be distributed to the various correlation centers, which should later upload their products to the main correlator.

4.3 Pipeline for Geodetic Post-Processing

Currently, there are no standard procedures for geodetic post-processing, although the IVS correlators all use the HOPS package for producing the databases holding the VLBI observables. There are a number of differences in terms of data handling, which may have an effect on the produced results and should therefore be investigated further.

An important aspect is the clock to be used for correlation. In Bonn we select a few scans both at the beginning and at the end of a session to determine the residual delay and delay rate of each station with respect to a reference station, which is usually Wettzell. At the USNO Washington correlator (WACO), the standard practice was until lately to choose a scan in the middle of the session instead. As a consequence the reference time of the clock model was different for IVS-R1 sessions, correlated in Bonn, and IVS-R4 sessions, correlated at USNO.

Other relevant aspects concern the post-processing of the data. Analysts should define best-practice guidelines for dealing with problematic data and, in particular, for channel flagging (the data are often affected by interferences) and for phase cal corrections (additive phases, manual pcal). The handling of Kokee Park phase-cal is still inconsistent between correlators.

4.4 Conclusions

The Bonn correlator is ready for the VGOS era, as has been demonstrated by the successful routine correlation of astronomical experiments with comparable data rates (e.g., the EHT at 32 Gbps).

The real challenge for VGOS will not be the computing power, but rather the data logistics. It is unlikely that stations and correlators will have at disposal adequate internet connections for e-transferring the amount of raw data generated in a VGOS session. Stations will likely have to ship their modules to the correlators, which is rather expensive. Furthermore, the foreseen duty cycle (24 hours per day on consecutive days) will require a rich media pool to provide stations with enough modules to keep observing while part of the raw data is being sent to the correlators.

References

[Deller et al. 2011] 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

[Bruni et al. 2014] G. Bruni, J. Anderson, W. Alef, A. Lobanov, & A. J. Zensus. Space-VLBI with RadioAstron: new correlator capabilities at MPIfR, in Proceedings of the 12th European VLBI Network Symposium and Users Meeting, Cagliari, Italy, 2014, ID: 119

[Gipson 2010] J. Gipson An Introduction to Sked, in Proceedings of the 6th IVS General Meeting, Hobart, TAS, Australia, NASA/CP-2010-215864, pp. 77-84, 2010.

[Bolotin et al. 2016] S. Bolotin, K. Baver, J. Gipson, D. Gordon, & D. MacMillan Transition to the vgosDB Format, in Proceedings of the 9th IVS General Meeting, Johannesburg, South Africa, NASA/CP-2016-219016, pp. 222-224, 2016.