The Bonn Correlation Center

Laura La Porta [1,2], Walter Alef [3], Simone Bernhart [1,2], Mikhail Lisakov [3], A. Müskens [4], Y. Pidopryhora [3], Helge Rottmann [3], A. Roy [3], Torben Schüler [2], Jan Wagner [3]

Abstract We present a status report of the Bonn Correlator for the years 2017 and 2018. 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 activities, focusing on aspects relevant for geodesy.

1 General Information

The Bonn Correlator is operated jointly by the Max Planck Institute for Radio Astronomy (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 is 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 Rheinischen Friedrich-Wilhelms Universität Bonn

Bonn Correlator

IVS 2017+2018 Biennial Report

2 Correlator Capabilities

Several versions of the Distributed FX software correlator (Deller et al. 2011) are available at the Bonn correlator, and in particular, a branch version developed by J. Anderson specifically for RadioAstron experiments (Bruni et al. 2014). For geodetic production we use the latest stable DiFX release, and before switching to a newer DiFX version we perform a comparison of the resulting observables. During 2017 and 2018 we have been using DiFX 2.5.2. 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 nowadays 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 to execute 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;
- 1.5 PB of disk space organized in RAID units (each with redundancy) and combined in a BeeGFS parallel cluster file system;
- 15 Mark 5 playback units; and
- nine 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 fixed disk arrays. 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 were already correlated in Bonn: Mk4, Mk5, DVP, and various components of VDIF.

The correlated data can be exported to FITS and HOPS (MK4) format. For post-processing the following software packages are available: AIPS, PIMA, and HOPS, 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 correlated in Bonn. They are the frontends to a local data-base into which all relevant information (observation date, participating stations, modules, and the status of the experiment) is collected.

3 Staff at the Bonn Correlator

The **MPIfR staff at the Bonn correlator** is a subgroup of the VLBI Technical Department, headed by W. Alef. Its members are:

W. Alef - head of the VLBI technical department, computer systems and cluster administration, manager of the BRoad bAND (BRAND) receiver project, VLBI expert, and consultant to the EU-VGOS project.

A. Roy - project manager for VLBI at the Atacama Pathfinder EXperiment (APEX), for DBBC3 commissioning, and head of the polarization conversion effort for Atacama Large Millimeter/submillimeter Array (ALMA) VLBI.

H. Rottmann - responsible for the beamforming software of ALMA, cluster administrator, DiFX developer.

G. Tuccari - guest scientist from INAF, leader of the Digital Base Band Converters (DBBCs) and the Fila10G development, as well as project engineer of the BRAND receiver.

J. Wagner - developer of DiFX, Mark 5 and Mark 6 software, responsible for correlation of EHT VLBI experiments, support scientist with instrumentation and observing-related work in mm-VLBI (EHT and GMVA).

Y. Pidopryhora - organizes, conducts, correlates and performs the post-processing of Global MM VLBI Array (GMVA) sessions and of various soft/hardware related tests.

M. Lisakov - joined the team in June 2018 to take care of the correlation and post-processing of the RadioAstron imaging observations. He also participates in the DiFX software correlator improvement for the needs of space VLBI.

S. Dornbusch - developer of firmware and software for the DBBC3 backend, responsible for manteinance of software for the DBBC2 backend, test and verification of the DBBC2 and DBBC3, support for stations that use a DBBC2 or a DBBC3.

The **geodesy group at the Bonn correlator** has 2.3 FTEs.

A. Müskens - scheduler of various IVS sessions, namely of INT3, EURO, T2, and OHIG, which he generates with the SKED software.

S. Bernhart and **L. La Porta** - coordinate the data logistics, prepare and supervise the correlation, carry out the post-processing and deliver the resulting observables to the IVS repository in form of databases. Besides these standard duties, they provide the stations with a feedback on their performance and support tests of the VLBI systems, in particular for the Wettzell Observatory.

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. The group members are involved in several astronomical projects, which are focused on very high resolution imaging especially with the Event Horizon Telescope (EHT).

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 contributes to the debugging of those tools.

4 Activities during the Past Two Years

IVS regular sessions: During 2017 and 2018, we correlated 97 R1, 12 EURO, 13 T2, 93 INT3, ten OHIG, and five days of CONT17. Since May 2018, the databases are produced solely via the vgosDBmake software in VGOS format.

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 could rely 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 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 2–3 weeks latency time. The final correlation of CONT17 was completed during February, but we did not resume standard activities until the beginning of May, due to some long-lasting discussions concerning the global set of clock parameters.

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 had to be reset often).

Tests of 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 testing of 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 a regular IVS-R1 session, i.e., R1.840.

Bonn acted as main correlator, therefore 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. No differences were found between the products of the various correlators, as was to be expected. We generated a new database for R1.840 by combining the output of the main and the branch correlators.

R. Haas performed a geodetic analysis of that database and compared the outcome with that of the original database, which contains only the outputs of the Bonn Correlator. Also this test was positive.

As a remark, 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 for producing the final database.

EU-VGOS Project: In March 2018—on initiative of W. Alef of the Bonn Correlator Center—started a collaboration with the three European stations of Wettzell, Onsala and Yebes, equipped with both standard S/X and VGOS systems, to carry out a VGOS Proof-of-Concept study. The aim of the project is to verify the processing chain for VGOS experiments end-to-end, from the scheduling to the geodetic analysis of the derived observables.

All parties are learning about the various aspects of the project, from the more technical ones, as system settings and data recording at the station, to data decoding and correlation, and finally to the post-processing of the data, i.e., to the fringe-fitting.

This is desired by all partners in light of the forthcoming IVS-VGOS sessions. In particular, most of the European stations have different back-end systems w.r.t. the American sites; therefore, the European stations must rely mostly on their own resources to debug their systems together with the correlator and the DBBC team in Bonn.

For example, these test sessions revealed that the VDIF multi-threaded files generated via VGOS broadband systems sometimes show anomalies. Namely, the threads are not correctly interleaved in time: for each time tag there should be all threads present in the file, whereas some threads appear with a certain delay w.r.t. the others. As a result, the DiFX's decoding buffers can overflow, thus leading to data loss. The percentage of correlated data decreased by 70% in the worst cases we experienced. The problem can be circumvented by reordering the threads in the raw data file and merging them into a single-thread via a DiFX utility called *vmux*.

RadioAstron: Approximately five experiments were correlated in Bonn in the second half of 2018. Those sessions involve up to 38 antennas and baseline lengths of several Earth diameters. The H-maser onboard the satellite stopped working in September 2017. Since then, there were two modes employed: a

Closed-loop mode (a.k.a. Coherent mode) as default and a Rubidium clock mode.

Global Millimeter VLBI Array (GMVA): Four sessions with up to 21 antennas were correlated in Bonn during the past two years. The recorded data rate is 16 Gbps for ALMA and 2 Gbps for the other telescopes, so that the amount of stored data can be as large as 700 TB.

As ALMA uses a different sampling rate, correlation requires intense use of so-called zoom bands, which allow to cut the non-matching observing sub-bands into pieces. These pieces are then stitched together so that the correct sub-bands can be reconstructed. Starting from Spring 2018, the GMVA network includes also the brand new Greenland telescope (GLT).

Event Horizon Telescope (EHT): The Bonn cluster is used also to correlate one half of EHT mm-VLBI experiments. The other half of the data is correlated at the MIT Haystack Observatory. The observing campaign of April 2017 lead to the first image of a Black Hole (The EHT Collaboration 2019).

The EHT campaigns in April 2017 and in April 2018 were carried out for five days using the phased ALMA and SMA, and up to seven single mm-VLBI antennas. The frequency setup consisted of multiple 2,048-MHz-wide IFs sampled by R2DBE backends. Each IF was recorded in dual polarization on separate Mark 6 units.

The April 2017 session had two IFs (2×2048 MHz dual polarization) and was recorded on two Mark 6 and eight modules at a total data rate of 32 Gbps. The April 2018 session had four IFs (16 GHz) and was recorded on four Mark 6 and 16 modules at 64 Gbps total. Aggregate rates are reaching 0.5 Tbps, with total storage requirements of around 5–10 PB for raw recordings, and 5 TB for the correlated and final polarization converted visibility data.

Correlation is limited by the available playback units; 32 Mark 6 with 32 expansion chassis would be required for a full 4-IF correlation assuming that eight stations participated in the observations. Hence, the correlation load is shared between the MIT Haystack and the Bonn MPIfR/BKG correlators. The full 230 GHz (1.3 mm) session is split by IF such that the Mark 6 modules of one IF subset are processed in Bonn and the other at MIT Haystack. Playback rate alone via fuseMk6 from a 2×2-module group is slightly above real-time and averages 2.4 GB/s total (18 Gbps).

The correlation of one hour of data for eight antennas in 1-IF takes about 3.5 hours. Format conversion, polarization conversion, fringe-fitting, and deliverables packaging via the pipeline of the EHTC Correlation WG need about the same amount of time.

Upon feedback from EHTC Calibration & Error Analysis WG a number of fine-tuned recorrelations have been necessary. For four IFs that led to a correlation and post-processing time of well over 20 weeks shared by the Haystack and Bonn correlators.

Digital Backends: Noteworthy results from 2017 and 2018 (Tuccari et al. 2018): acceptance and use of the DBBC3 in DDC mode for VGOS observations and in OCT mode for EHT observations; further improvement of DBBC2 capabilities in terms of sampled bandwidth and performance.

BRAND: During the past two years the main tasks have been: a) a survey of the specifications which would enable the installation of a BRAND receiver at the EVN antennas; b) definition of the receiver frontend with the primary focus feed including filters for RFI suppression; c) selection of suitable chips for the sampling board of the backend; d) development of a high data rate processing board with FPGA processors. This board has an input data rate of almost 1 Tbps and up to 512 Gbps as output and is the most powerful data processors available in radio astronomy at present.

5 Current Status and Future Plans

All activities described above are ongoing and some more have been started in 2019.

- M. Lisakov and J. Wagner will implement a closed-loop mode correlation in DiFX to be able to deal with the RadioAstron sessions observed after September 2017. On 10 Jan 2019 the satellite stopped to receive commands at all. Since then, no observations were performed, except for GG085F, which was observed in spring 2019 solely with the Global VLBI. Thirteen experiments await correlation to complete the project.
- Recently the upgraded NOrthern Extended Millimeter Array (NOEMA), the next generation of

IRAM'S Plateau de Bure instrument, joined the GMVA network. A first fringe test was successful.

- During the GMVA session carried out in Spring 2019 the VLBA recorded for the first time on Mark 6 units so that the recording data-rate for the GMVA can be increased now to 4 Gbps.
- Lately, some preliminary tests have successfully been performed at Pico Veleta by using DBBC3s.
- J. Wagner is working on DiFX outputbands feature for DiFX 2.7. The feature is based on earlier standalone de-zooming 'difx2difx' he developed to recover and correlate GMVA+ALMA 2017 (Issaoun et al. 2019). It will enable correlation of VLBI experiments that have inconvenient sky-frequency IF placements and bandwidth mixtures of (non-)overlapping 32/58/62.5/64/128/2048 MHz wide bands.
- He is also working on a Mark 6 reading layer in DiFX together with FUSE layer fuseMk6, as well as improvements towards multi-datastream handling in DiFX native Mark 6 support. He implemented some improvements for playback of lossy or "clumpy" multi-threaded VDIF recordings affecting some VGOS stations and for the PolConverter conversion flow.
- A. Müskens started testing the VieVS (Madzak et al. 2013) scheduling program in collaboration with the Vienna group. In 2019 several INT3 sessions were scheduled with the VieVS software with good results. In early 2020 A. Müskens will retire, thus ending the IVS scheduling activity in Bonn.
- In September 2019, W. Alef will retire and his role as group leader will be taken over by H. Rottmann.
- In August 2019, L. La Porta will also leave her job at the correlator. In April a new scientist was employed by the Reichert GmbH to take care of geodetic correlation with S. Bernhart: Y. Choi.

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

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 their disposal adequate Internet connections and data storage 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

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

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

G. Giovannini, T. Savolainen, M. Orienti, M. Nakamura, H. Nagai, M. Kino, M. Giroletti, K. Hada, G. Bruni, Y. Kovalev, J.M. Anderson, F. D'Ammando, J. Hodgson, M. Honma, T. Krichbaum, S.-S. Lee, R. Lico, M. M. Lisakov, A.P. Lobanov, L. Petrov, B. W. Sohn, K. V. Sokolovsky, P. A. Voitsik, J. A. Zensus, and S. Tingay A wide and collimated radio jet in 3C84 on the scale of a few hundred gravitational radii Nature Astronomy, 2018, Vol. 2, 472-477, http://adsabs.harvard.edu/abs/2018NatAs...2..472G

G. Tuccari, W. Alef, S. Dornbusch, M. Wunderlich, A. Roy, H. Rottmann, J. Wagner, R. Haas, K.-A.L. Johansson DBBC3- the new wide band backend for VLBI in Proceeding of the 14th European VLBI Network Symposium & Users Meeting (EVN2018), 8-11 October 2018, Granada, Spain, PoS (EVN2018) 140

S. Issaoun, M.D. Johnson, L. Blackburn, C.D. Brinkerink, M. Mocibrodzka, A. Chael, C. Goddi, I. Mart-Vidal, J. Wagner, S. Doeleman, H. Falcke, T.P. Krichbaum, K. Akiyama, U. Bach, K.L. Bouman, G.C. Bower, A. Broderick, I. Cho, G. Crew, J. Dexter, V. Fish, R. Gold, J.L. Gmez, K. Hada, A. Hernndez-Gmez, M. Janen, M. Kino, M. Kramer, L. Loinard, R.-S. Lu, S. Markoff, D.P. Marrone, L.D. Matthews, J.M. Moran, C. Mller, F. Roelofs, E. Ros, H. Rottmann, S. Sanchez, R.P.J. Tilanus, P. de Vicente, M. Wielgus, J. A. Zensus, G.-Y. Zhao APJ, Vol. 871, Issue 1, article id. 30, 17 pp. (2019)

The Event Horizon Telescope Collaboration First M87 Event Horizon Telescope Results. II. Array and Instrumentation APJL, Vol 875, Issue1, article id. L2, 28 pp. (2019)

M. Madzak, S. Böhm, H. Krásná, and L. Plank. Vienna VLBI Software Version 2.1 User Manual web document at http://vievs.hg.tuwien.ac.at/vievsDoc.pdf