Max-Planck-Institut für Radioastronomie

Bundesamt für Kartographie und Geodäsie

# The Bonn Correlator and VGOS sessions

Laura La Porta

on behalf of the VLBI Correlator team in Bonn

W. Alef, H. Rottmann, S. Bernhart, L. La Porta, A. Müskens, T. Schüler, J. Wagner

Max-Planck-Institut für Radioastronomie

# The Bonn Correlator

Bundesamt für Kartographie und Geodäsie

**MPIfR** → infrastructure (host institute) + hardware + connectivity + staff

**BKG** → hardware + *outsourced* staff (Reichert GmbH)

**IGG** → internet connection + staff

- **DiFX software correlator** - latest stable version frozen for correlation production in geodesy DiFX 2.5.2
- **High Performance Cluster (HPC)** replaced in 2015 to match both VGOS and mm-VLBI requirements

Max-Planck-Institut für Radioastronomie

# The Bonn HPC

Bundesamt für Kartographie und Geodäsie

- 68 nodes x 20 compute cores = 1360 cores → 10 x higher computing power w.r.t. old cluster
- 3 head nodes → possible to run more correlation in parallel
- 56 Gbps Infiniband interconnect between nodes
- Storage space > 1 PB, organized in BeeGFS file system

Max-Planck-Institut für Radioastronomie

# The Bonn HPC

Bundesamt für Kartographie und Geodäsie

- 15 Mark-5 playback units
- 7 Mark-6 units with 4 bays

- 2 x 1 Gbps Internet connection -

  2 Provider with fiberglass nearby +
  cheaper connection with time
  → more bandwith expected soon!

MPIfR staff is responsible for
- cluster software upgrades
- cluster hardware maintenance and repair
- general IT-support
- software correlator improvements (DiFX developers)

Max-Planck-Institut für Radioastronomie

# MPIfR Staff @Bonn

Bundesamt für Kartographie und Geodäsie

- MPIfR VLBI technical department headed by W. Alef:

  H. Rottmann, A. Roy, J. Wagner, Y. Pidopryhora, M. Lisakov, G. Tuccari, S. Dornbusch

  R. Märtens, H. Sturm, M. Wunderlich

- Projects

  Implement & upgrade VLBI @ APEX (Atacama Pathfinder Experiment) for EHT

  DBBC2 & DBBC3 backend development

  BRoad bAND broadband receiver development

  Atacama Large Millimeter/submillimeter Array (ALMA) Phasing Project

  RadioAstron correlation and data transfer to ASC

  EU-VGOS

  DiFX developers

  Support of VLBI operations (upgrades for EHT and GMVA) at Pico Veleta and NOEMA

Max-Planck-Institut für Radioastronomie

# Geodetic Group @ Bonn

Bundesamt für Kartographie und Geodäsie

- A. Müskens (IGG), S. Bernhart and L. La Porta (Reichert GmbH)

Scheduling (SKED) of INT3, T2, EURO and OHIG IVS-sessions

Mainteinance of the webpage for e-transfers to correlators

Preparation and supervision of correlation

Post-processing (fringe fitting) – HOPS

Database submission to IVS repository – vgosDBMake

Feedback to stations

…..

Correlated in 2018

52 R1s
6 EUROs
7 T2s
49 INT3s
6 OHIGs

Max-Planck-Institut für Radioastronomie

# EU-VGOS test sessions

Bundesamt für Kartographie und Geodäsie

dual linear polarization (X,Y)
4 x 512 MHz bands
8 frequency channels per band
2-bit sampling

acquisition rate of 8 Gbps

4 hrs observations (~150 scans)
3 stations
geodetic schedule (scan duration = 30 sec)
+ calibrators (120 sec)

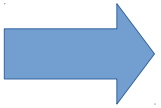

~6 TB per station

station – Bonn correlator: 600 Mbps connection ==> e-transfer nominal duration

1.5 – 4.5 d

data transfers running in // (best case scenario) - start transfers one after the other

In practise – VGOS not yet in production phase in Bonn – 2-4 weeks to fetch data

Max-Planck-Institut für Radioastronomie

# EU-VGOS test sessions

Bundesamt für Kartographie und Geodäsie

- Conversion of multi-threads into single-thread data (file based correlation issues)
  → 0.5 day for 6TB
- Correlation time (128 spectral channels, 3 stations, ~150 scans of 30 sec lenght)
  → 2-3 hrs
  (correlator needs in average 1-1.5 minutes to complete one job)
- Post-processing (fringe-fitting + DB creation) → ??? (1-2 days)

**Delivery of data** to collaborators:

best case scenario (reserved bandwidth for transfers) — after 1.5 d + 0.5 d + 2.5 hrs + 1 d ~ **3.5 d**

(conversion runs parallel to transfers) — 5 d + 3 hrs + 4 hrs ~ **6 d**

In practise – VGOS not yet in production phase in Bonn – 1 month to fetch data

Max-Planck-Institut für Radioastronomie

# IVS-VGOS intensives

Bundesamt für Kartographie und Geodäsie

1) 2 participating stations: Kk, Ws

2) 8 Gbps,

3) 1 hr session,

4) 15 sec on source, 30 sec slewing for ~ 80 scans / session → 15GB data/scan

5) 800 Mbps e-transfer rate for both station and correlator

6) correlation time

7) post processing & DB generation

1.2 TB / station

Transfer time / session ~ 200 min

Correlation time ~ 1 hr

Post-processing time ~ 5 hrs ?

Delivery within one working day, if observations starts at UT 0700

Max-Planck-Institut für Radioastronomie

# Future VGOS sessions

Bundesamt für Kartographie und Geodäsie

How to extrapolate these latency time to sessions with other parameters (nr of stations, acquisition rate, duration … )

Possible limiting factors to DiFX performance

| | *in Bonn* |
|---|---|
| ✔ Read I/O perfomance (datastream nodes) | 16 Gbps |
| ✔ playback rate: | |
| ✔ Mk6 units – tested successfully 4-6 Gbps, via J.W.'s fuse mounting | 56 Gbps |
| ✔ computing capabilities | 1360 cores |
| ✔ storage capacity | 1.3 PB |
| ✔ write I/O (head node) | (not relevant in geodetic case) |

Computing time increases linearly with amount of data to be correlated on a High Performance Cluster with the same characteristics as in Bonn (Rottmann – priv. comm.)

Max-Planck-Institut für Radioastronomie

# VGOS 24 hrs-sessions

Bundesamt für Kartographie und Geodäsie

IVS-VGOS
24 hrs sessions 2019

8 Gbps,
6 sites,
24 hrs sessions,
580 scans (30 s on-source)

→ ~16-17 TB per station

- e-transfers → 4.25 d/st, max 3 transfers at 600 Mbps in parallel → 10.5 d minimum
- 2 modules per station (possibly only 1 - Yj experience) → 1-2 week to receive all modules, but shipping costs!
- correlation time → 10-15 hrs
- post-processing time (fringe-fitting with HOPS + DB creation) → 3-5 days

**Latency time**: 10.5 d + 5 d (correlation + post-processing) ~ **11-12 d** best case scenario
Haystack experience ~ **1 month** (mostly due to shipment)

Nowadays bottle neck is the transfer of the data from stations to correlator!

Max-Planck-Institut für Radioastronomie

# Open issues

Bundesamt für Kartographie und Geodäsie

- Computing time could be significantly improved by leaving the scan-based correlation scheme behind
- Post-processing ???
- Main issues for VGOS are the transfer and storage of the data
- Currently the internet connection costs are expensive (situation may vary strongly with Country)
  Stations need 2 Gbps connection and correlator 10 Gbps to decrease e-transfer time
- Station should be equipped with local buffering systems (if using flexbuff for recording)
- Nr of Mk6 modules at stations = 2 x modules needed for sessions in one week (latency time)
- Storage at correlator (paid by stations?)
- Shipping costs

Max-Planck-Institut für Radioastronomie

# MPIfR Staff@Bonn Corr.

Bundesamt für Kartographie und Geodäsie

Astronomical correlation

- handled by MPIfR staff and support scientists
- focused on *very high resolution* astronomy

  **Radioastron** – 15 sessions with up to 20 antennas (baseline lenghts of various Earth diameters)

  **Mm-VLBI** – 2 Global 3mm-VLBI Array (GMVA) sessions per year with up to 15 antennas
  2 Gbps data-rate (4 Gbps with ALMA)
  up to 700 TB raw data

  Event Horizon Telescope (**EHT**) @230 GHz half session correlated in Bonn

- include also tests for development of DBBC VLBI backends

Max-Planck-Institut für Radioastronomie

Bundesamt für Kartographie und Geodäsie

# Distributed Correlation

- Aim at testing distributed correlation for future VGOS sessions
- Each correlator only receives the raw data for part of the session.
- Main correlator → fringe search, preparation of vex and v2d files for correlation, post-processing and database creation
- Branch correlators → correlate and upload correlated data to main correlator
- Downsides:

  More complex logistics – stations must transfer the raw data to various correlators
  correlators must upload correlation results to main correlator

  Main correlator can not easily correct for clock jumps –
  possibly need to fetch more data in the middle
  based on log Information or on station‘s start/stop message

Max-Planck-Institut für Radioastronomie

Bundesamt für Kartographie und Geodäsie

# Distributed Correlation

R1840

- Main correlator: Bonn → vex, v2d file and control_file for fringe fitting
- Branch correlators (1 hour of data): Onsala (R. Haas)
  Warkworth (S. Weston)
  Hobart (J. McCullum)
  Vienna (J. Grüber)
  Seshan (F. Shu)
- DiFX version 2.5.2 (correlation,conversion to Mk4 format), HOPS 3.18 (fringe fitting)
- No significant differencies were found when comparing both the difx outputs and the converted mk4 data
- R. Haas will analyse the R1840 vgos databases
  which contain solely the Bonn correlated data
  by merging the difx outputs (or Mk4 data) of the various correlators

70 Mbps per thread processing time
15 GByte per scan => 120000 Mbit / 70 Mbps = 1714 sec / thread

Processing with 10 nodes x 20 cores = 200 threads => ~ 8.5 seconds/scan

Pure computing time for 80 scans => 685 s

About 5 s overheads for starting and stopping each job => 80 x 5 = 400 s

Total processing time for one job => ~ 1000 s = 16 min

Based on our experience, on average: processing time / job ~ 1-1.5 min

Correlation time is not an issue!

Still room for improvement – scan based correlation not computationally efficient due to overheads !!

Max-Planck-Institut für Radioastronomie

# The Bonn Correlator

Bundesamt für Kartographie und Geodäsie

- Raw data are currently recorded at the stations
  - on Mark-5 or Mark-6 modules (only astronomy for the moment)
  - on flexbuff

Modules may be shipped to Bonn, or the data are directly e-transferred to the HPC (mostly the case for geodesy) via two 2-Gbit connections

- Various data formats have already been correlated in Bonn: Mk4, Mk5, VDIF, DVP and Mk6

- Correlator outputs and other important files (e.g. Mk4 format directories, vex and v2d files) are backed-up daily on the HPC.The final products are also archived on the MPIfR archive server (will be kept for 10 years).

- Data export is possible in FITS and HOPS (Mk4) format. Post-processing is done with HOPS, AIPS and PIMA.

- EXPAD/COMEDIA have been expanded to help bookeeping the experiments correlated in Bonn by collecting all relevant informations (observation date, participating stations, modules, status of the experiment).

Max-Planck-Institut für Radioastronomie

# Bonn as IVS Correlator

Bundesamt für Kartographie und Geodäsie

**Need for**

**a standard pipeline for geodetic correlation and post-processing**

- Fringe search – A priori clock and a clock_rate are estimated for each station by performing a linear fit of the logged GPS-fmout values.

We pick some good scans at the beginning and at the end of a session to correct clock and clock-rate values w.r.t. a reference station

- We always apply additive phases before database production. That is often not the case for other correlators.

- Correlator report content

**Guidelines on how to handle various cases during post-processing**

Analysts should agree on what are best practise in various cases:

- Channel flagging (also baseline based?)
- Manual pcal (always in both bands?)
- Notch filters