

ZWURM, 29-06-2020 14:00 (WURM through Zoom because of #COVID19 house quarantaine/wk17)

Present eBob, Mark, Des, Harro

Two absentees due to EAS special session, one verloffing.

Des: Working on last RINGS deliverable. Following finishing that will work on building CASA packages for MichaelJ; has been in contact with Ville from NRAO because building CASA5 packages out of CASA6 tree only seems to work locally. Mark/Des will contact MichaelJ to see if he can test on casadev (account should still be available). Will also interest MichaelJ to test the wide-band fringe-fitter. Following that will work on the stacking correlators functionality (stacking polarizations which are called correlations in CASA, red.).

eBob: Found the experiment having e-Merlin > 1 datastreams. Copied to other flexbuff and renamed, for testing. Needed some changes in the DB (test-db used now), requires extra permissions. During last week's e-VLBI TianMa entered late in the experiment and operators used simultaneous clocksearch + prod correlation feature but forgot to use the correct datasource -> results in lot of issues. Will write short HowTo for the more esoteric e-VLBI modes. Experienced e-transfer problems with e-Merlin; one out of their three flexbuffs had file system issue causing the e-transfer to hang, waiting for a reply that didn't come. Fixed by adding a timeout on the command (2h); informed e-Merlians - fb now rebooted/fixed. The RadioAstron disk pack recovery succeeded - no disks seem to have been faulty so why Mark5 playback didn't work anymore = anyone's guess. Since Mk5 is dead, will not investigate.

Mark: Finalizing gaincurve code; GeorgeM looking into how to integrate in gencal, to not have to write a new approach. Working on getting the FITS-IDI gaincurve table into CASA table into casacore. Investigating FT036 (8 Gbps) some more: phase-cal extraction is on, inefficient implementation, causes some slow-down. Now up to 2 Gbps (4x slow-down; 5m scan takes 20m, used to be ~1hr!). Maximum measured flexbuff input node performance so far is 3.5 Gbps so can go only two times faster but not clear where next bottleneck is. Two stations on one flexbuff not the issue; tested by removing 2nd station from correlation - no speed-up; disk I/O and network not bottleneck so far. Picking up V0 work, finding out where left off.

After the ZWURM discussed the backup situation in light of the LT08 arrival - aggregation of three backup systems onto one physical tape. Need to keep user interface simple and compatible with current scripts. eBob will start to investigate current scripts and how we could merge them into non-local backup/restore. Need to contact PaulB about burp (backup system) requirements for tape backend.

Also discussed polconvert-on-MS/reduction of FITS-IDI size by factor 1/2 by storing only parallel pol after polConvert [driver: one experiment, five epochs, without such sizing fix, will double the

total EVN archive size - i.e. all EVN data since 1998 or so]. Most likely current approach: create one FITS-IDI file, run polConvert, generate Jones matrix from that calibration, write script to apply to all other MSs, update tConvert to allow writing subset of polarizations.