

ZWURM, 22-06-2020 14:00 (WURM through Zoom because of #COVID19 house quarantine/wk16)

Present eBob, Mark, Paul, Ilse, Des, Aard, Harro

General announcements:

Ilse sais that ASTRON has decided no events take place at ASTRON until the end of the year. For the CASA VLBI workshop now investigating holding it at De Börken in Lhee.

eBob: datastreams in diskrecording code now done, not tested; how to build a test case [maybe e-Merlin multi-file recording still present on flexbuff]. MartinL reports problem correlating RadioAstron data - disk pack record pointer reset to 0. All recover commands failed. This is a PATA disk pack so cannot easily load disks in e.g. Mark6 (SATA); now copying disks one-by-one to flexbuff to be able to recover disk pack contents. Still mystery how the reset came about: until now hypothesis was this is due to mismatch between firmware used to erase the disk and the one recording it. Disk pack played back fine earlier so the reset had to have happened whilst at JIVE. Working on the merge of SCHED 11.6 into pySCHED, mostly catalog work.

Ilse: last-minute EAS things ramping up. The flowchart for the pipeline is shaping up nicely. Attempting to get AOFlogger working in the notebook revealed some issues - André0 is writing v3, so maybe target that version in stead of spending a lot of time on the current version. This week and next swamped with meetings - everybody wants to have a last meeting before summer.

Aard: investigating GA043A crashes reported last week. Caught SIG13 in the code; delay crash, enough to dump state. Seems fb14 (and upon closer inspection: fb13 and fb3) are always in when crash i.e. not necessarily fb14 root cause. Seen instances where corr node continues whilst input node holds. Found that systemcall wrapper returns size\_t (unsigned) but returns the literal return value of write(2) systemcall (signed!) - upon error negative value gets returned as very very large positive value. Fixed but does not solve issue. Question remains: why does the write(2) call return an error - i.e. why does the TCP connection become broken? AOFlogger only has Py3 bindings which do not easily work in CASA5 which is Py2 based. AOFlogger v3 (being written) uses pybind11 and can generate Py2/Py3 bindings, so probably wait/use that AOFlogger in Jupyter CASA image. WSClean does not have Python bindings at all so probably resort to os.system or subprocess.Popen in Python to drive from notebook. FranzK reported issues with SFXC filterbank output: polarizations written in wrong order and request for t\_int < 1 us, which is current SFXC limit. There should be no obvious reason for this limit anymore but will require a fair amount of work so RoI might be small. Also before/after pulse profile small dips in spectrum observed, compared to digifill (from PSRCHIVE). Probably related to not per-FFT van Vleck correction (which is what digifill does). Interesting discussion to be continued outside (Z)WURM.

Paul: Invited to be on a CERN WhiteRabbit improvement panel discussion this Friday. Set up a Mattermost team for CRAF. Wrote letter of support for French time- and frequency transfer project. Dealing with failing disks. Working on monitoring disk usage using cacti on graphs.jive.nl; fair bit of work: default SNMP counters = 32 bit = 2 TB file system size limit. Now made 64 bit counters himself to monitor. Received a certificate for code.jive.eu. Reported a number of bugs in SURFnet network dashboard. Working with KlaasS on ASTRON 100 Gbps migration now that the h/w is there.

Des: looking at wide-band fringe fitter and is happy: works for NMEs fine. Will write document and verify it works for blackhole cam data. MichaelJ did report a problem (with test case sent), will investigate. Database was prepared for VEX2 operations but db2vex/vex2db weren't; now working on BobC reported issues. The last RINGS deliverable is nearly done, C++ dispersive fringe fitter works really nicely on VGOS data, now writing up.

Mark: gaincurve code works. Wrote a tool to add gaincurve to FITS files (should move towards using this at JIVE). Tested on EVN experiment: using real GC in stead of interpolated sampled yields diffs of 0.2/0.3%. Also tested on VLBA w/ non-trivial GCs, found that VLBA reports as function of zenith angle; EVN records as function of elevation. The script now deals with both correctly. ESCAPE WP3 is planning a virtual workshop in/really just before summer on software lifecycle, versioning, continuous integration, licenses and management, four/five days @ max three hours/day. JIVE will contribute a bit, seeking contributions of other partners.