

ZWURM, 30-03-2020 14:00 (WURM through Zoom because of #COVID19 house quarantine/wk4)

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

Mark: Most time spent on CASA VLBI requirements document. Got a lot of response, feedback already incorporated. Efficiency of working from home < 100%.

Aard: Investigated m-node problem of last week. Follows m-nodes (shuffling ranks/machine files) so related to 18.04 (other nodes are 16.04). Tried MPI upgrade 4.0.1 -> 4.0.3, no difference. MPI source code bug: reports EINPROGRESS in stead of real error message (ETIMEDOUT). Experiment with tcp syn timeout to very large value: seems to fix issue but is a hack. At startup out.sfxc alloc ~110GB of memory -> machine very busy, probable cause? Now staggered startup of MPI processes (60us between), adds ~ 1 minute to startup time, no more issues. Will experiment with different sleep intervals. Correlating one scan could take several hours, cause: channels with + without e-Merlins. Channels w/o have 1/2 baselines = 1/4 less work => nodes are waiting for more work, so after a while correlator runs with single-core performance. Fixed by increasing slice time from 0.25s to 1s. The watchdog timer (30min) triggered several times (correlation takes 80-90min) so something's not right there. [Mark: possibly enter into some deadlock and then times out] [Action harro: side meeting on correlator performance]. Finally success with Jupyterhub docker spawner! Behind scenes spawner generates new container (only visible with "docker ps -a") with part of configuration in, which is reused. Deleting that container and rerunning = now OK!

Ilse: Notes that with four children trying to follow online school and working, things are sometimes difficult to schedule. Ran new dataset through notebook without errors, but calibration results need verification for sensibility. Identified problems: the special scripts to be run before (should be available from inside notebook) and plotms still giving problems. Next up is a non-phase ref dataset to investigate whether a separate notebook is needed or not. Mark's CASA VLBI requirements document triggered ionospheric calibration ideas: what happened to James Anderson's work? [Mark: in ALBiUS conclusion was that w/o really good ionospheric model it's 50/50 between improving or worsening data]. Working on CASA workshop tutors list, EVN symposium (together with Olga). Waiting for EAS decision on how to deal with COVID-19, should be soon because tomorrow deadline for invited talks.

Paul: Working on m-node problem from network point of view: generated (missing) cacti graphs, did tuning and improvements, now sampling faster again whilst collecting more probes. Result wrt m-node problem: inconclusive from network PoV. Got an LT08 quote. Working with Klaas on ASTRON 100Gbps upgrade. Support requests (install paramiko, spawn new VM, DB mgmt) handled. Physical access to building will likely be postponed for a few weeks more: likely disk #FAIL or powerfail. How to service the former and bring

system(s) back up in case of the latter, if none of the techs have access to the building [action Harro].

Des: Worked on papers database, prio is on linking EVN project codes. Difficulty: can run full-text query ADS but not get access to full text. Ran into some fringe-fitter unit test program issue, need to really understand calibration framework.

eBob: Working on getting datastreams into production database, copied DAS table manually from test => production DB (Des' script, if any, may have to be aware of this). runjob starts OK but corr was busy so not tested beyond that. Currently back to production so will check with operators. Working on NorthStar GDPR checkbox implementation - can't connect to dop800 database, investigate with Paul after meeting. After consulting with supp.sci. pySCHED input not terminating with "/" will remain error. pySCHED install on eee-dev issue: procedure tries to compile pyqt from source, #FAIL miserably. In setup.py forced to previous version of pyqt, which installs from a wheel and then it works. [Mark: can wheel be forced? pip cmd line yes, setup.py no. Also not good to restrict on method, possibly restriction to known correct version is more important].

Harro: After jive5ab 3.0.0-rc1 got a number of compilation errors, caused by newer compilers on several different platforms, with more strict checks - revealed at least one proper bug. Working on jiveplot Py2+3 compat, can now run standardplots again. Would be nice to rename branch because current name is very not indicative of this.