

WURM, 17-04-2023 13:00 WURM (Library meeting room)

Present Bob, Aard, Wybren, Des, Mark, Ilse, Paul, Marjolein

Plenary part: Marjolein asks everyone to start writing tasks/projects they could be, are actively or waiting for someone/something on (coloured) post-its and put them on the white board in the respective category. This will provide an overview of all work our group is responsible for and allow us to see which projects should/could be prioritized or finished and at the very least give everyone an idea of the workload. It's probably not the best solution in the world but let's see where this leads.

Ilse: NOVA ISC in Leiden big part of last week; rest spent on NAEIC mails & work climate survey (which likely won't happen before NAC) and CASA workshop organizational business. EHT review nearly done. This week start on ORP EVN proposal and start preparing workshop lecture.

Aard: clocksearch VEX parser turned out to not be VEX2 proof (fixed), some fields changed location. Checked GMRT vdif data: several issues detected (in one file all samples having the same value, data rate in file not matching expected setup). Q: vdif thread ids in the file - should the s/w trust them? [Mark/Marjo: yes, that's why the IDs are explicitly recorded in the VEX file]. Implemented amplitude normalisation based on sampler stats, seems to fix long standing issue w/ mixed bandwidth correlations where each band was normalized individually; now commissioning this [Mark: there may be corner/edge cases, need to be aware of those] [Marjo: write commissioning proposal for ZsoltP/BobC to execute]. Experimented with toy correlator data ingestion - likely will do an offline cornerturn and then just load data into GPU [Mark: should coordinate internally the division of labour].

Bob: VEX2 issues reported by BobC (log2vex, pySCHED) coming; expect to do the work twice: 1st round log2vex hacks for immediate operational support, 2nd round is to move some of the work into pySCHED; fixed pySCHED printing: some fields were moved to a different location; found Out-of-Bounds write issue in the FORTRAN code if #-of-channels > 31, leading to erroneous mode name written to file (fix: limit printing to 31 channels). numpy PR response: maintainer seems ok with accepting, result in new release + possible backport. Experimenting with OAuth2 in BHTOM. Switching off Mark6s used for VLBA disk pack copying when not used side effect: monitoring tools poll and give errors, not only Zabbix btw.

Des: investigating/assessing Apache Arrow in preparation of EHT s/w and data compat wg meeting and RADIOBLOCKS: seems OK format but no certain statements by anyone on whether it will or won't work with Dask; EHT wg open to non-member participants. Resurrected closure phases in dask code, getting non-zeroes, turns out source has structure (good). Working on publishing first real DOI; related: working on tooling to get DOIs into archive - script expects all epochs present already. Making edits in db2vex for VEX2. PolConvert

in c++ runs w/o #FAIL but not sure it works.

Wybren: Zabbix community update installed, did not fix Ansible <-> Zabbix module issue(s). Moved mark5info to another machine, not all info visible? Moved Flask to systemd unit (for Des' app(s)): can now easily start/stop service(s); reverse proxy to them works, need to check that LDAP traffic does not go over unencrypted conn between proxy and actual server. Have four broken disks this week. Detected uptime < 10 min in Zabbix: investigation reveals uptime counter had wrapped b/c > 496 days or so. Verifying permissions of database accounts: scope to create specific accounts with less permissions. State of Py3 on out.sfxc? [Answer: no change]. Will repurpose cl[01] for proper Zabbix test.

Paul: archive2 copying data progress: now at 180 TB from 250 TB; should be done in ~5 days; still only getting ~1.2 Gbps - cause? maybe rsync, maybe source pool (very old disks on archive). If Mark6-0 used get pkt errors. Interesting SURF ticket: link Zwolle - Groningen break, "ship will arrive North Atlantic in ~four days". Submitted NAC abstract; working on WR paper, found bug in latest WR firmware.

Mark: IVOA RiG meeting, preparation for May interop: ObsCore visibility data extension still under discussion. EOP correction script produces cal table in MS, use applycal() in CASA; verified by recorrelating recorded e-VLBI w/ good EOPs and comparing to corrected previously correlated data; big red label: this produces MS with different phase convention; new j2ms2 adds keyword to MS, script tests for kw to apply appropriate amount of minus signs; handed script to supp sci to be run on data. EHT s/w + data compatibility mtng: julia is still (again?) a thing (new code being developed); data format: lcd = UVFITS, requires new one: producing memo w/ input from Mark, Des, looks like converging on "MS-like data model". Will bring it up in this week's CASA VLBI coord mtng. Handling CASA workshop finances.

Marjolein: Created pair of tConvert/j2ms2 linked against VEX1.5/2 compat library with local mods. EOP correction needs tConvert that's capable of exporting other data column, that feature is part of the "export subset of polarizations" branch which was merged with master that has other features integrated after that branch (any size FITS, lis-file support for tConvert) and linked to VEX1.5/VEX2/local mod library. Whilst at it implemented trick to compile actual/current git status into binaries; j2ms2 and tConvert then gain --version to show just that.