

WURM, 9-12-2019 14:00

Present Aard, eBob, Paul, Des, Harro

Aard: 3x SFXC freezes (no more output) in last week's e-VLBI even though very standard observing mode. 1x MPI tcp error + 2x unknown. In those it happens in the same experiment, each time when one station stops sending data. So far unsuccessful in reproducing locally (we do not have enough h/w to fully simulate 17 stations e-VLBI - requires 17 generators (=9 flexbuffs) + 17 input nodes (=another 9 flexbuffs). A new SFXC release should be done before next e-VLBI; "trunk" has many fixes that should be migrated to stable. Requires coordination with Mark. Been busy writing the follow up SFXC paper.

eBob: Communicated with e-Merlin about what they do with our hand-crafted VEX files: not much but look at the \$FREQ section, so now looking into automating e-Merlin schedule creation in pySCHED. Working on flexbuff_copy_manager to support data recorded on Mark6 disk packs. Updated database to reflect the fact that Arecibo now behaves as VLBA station wrt data format (VDIF 5000 byte, threading). VLBA Mark6 disk packs now copied using vbs_fs FUSE file system (only 200MB/s [suggestion: use "-n X" with 1 < X < 6 command line option to "vbs_fs"]).

Paul: Quick & dirty Mark6-1 install w/ recent O/S to get it going. Mark6's should get external IP address. Cluster migration is almost done (8 nodes to go), will require some more hardware to be ordered (some or all of eth cards, optics, fibers). A solution to the towers on a chessboard chess problem was used to automate flexbuff-flexbuff iperf tests. Results are intriguing: fb0 has packet loss sending to anyone, fb13,14,15 sending to everyone (fb15 has 25G optics), fb12 is OK despite being on same switch port as 13-15. Will conduct several tests to see if this can be linked to fiber, card or optics. The 100G upgrade won't happen this year because of KlaasS going on holiday soon and whom this cannot proceed without. Because of O/S upgrade the worldmap creation script does not work anymore and Benito c.s. want a bespoke worldmap for their upcoming paper.

Des: The C++ dispersive fringe fitter did not work on test case whilst Python version worked like a charm. Found to be differences in C++ convergence criterion and the scaling of numbers was different between C++/Python. When dispersion ~ 0 , takes a few iterations before the results start to change, updated implementation to run at least a few times. The unit test failed, most likely because it was not updated to be in sync with the new code producing different (allegedly better) results. Will be checked if indeed better + updated.

Harro: Post Mark6-1 O/S install a few administrative actions have to be taken to make it ready for copying data off. Documented in JRM ticket, should be automated via Ansible. Mike Sipior has provided us with a development EEE (virtual) machine where ESCAPE work can take place.