CORRELATOR OPERATIONS REPORT,  JIVE
EVN TOG MEETING,  June 2015, Robledo

19 June 2015  (statistics cover 27 Sep 2014 – 19 Jun 2015)
Bob Campbell

SCIENCE OPERATIONS

The table below summarizes projects correlated, distributed, and released from 27 September 2014 to 19 June 2015.  The table lists the number of experiments as well as the network hours and correlator hours for both user and test/NME experiments.  Here, correlator hours are the network hours multiplied by the number of multiple correlation passes required.  This definition carries over to the EVN software correlator at JIVE (SFXC), even though it may run faster or slower than real time.  Because of its enhanced capabilities, multiple correlator passes for SFXC occur only for phase-referenced spectral-line observations (separate "continuum" and "line" passes) and for pulsar observations wanting different gating/binning configurations.

| | User Experiments | | | Test & Network Monitoring | | |
|---|---|---|---|---|---|---|
| | N | Ntwk_hr | Corr_hr | N | Ntwk_hr | Corr_hr |
| Correlated | 62 | 549 | 681 | 18 | 51 | 51 |
| Distributed | 50 | 431 | 497 | 16 | 36 | 36 |
| Released | 48 | 438.5 | 469.5 | 21 | 51 | 51 |

(3 user experiments from sess.3/2014 and 1 from sess.1/2015 were abandoned by their PIs prior to correlation)

The following table summarizes by session the user experiments with activity since the previous report (entries = remaining to do / total).  The "(e)" and "(d)" refer to e-VLBI or disk experiments within a traditional EVN session.

| | N_to.corr | Corr.hrs | N_to.dist | |
|---|---|---|---|---|
| Sess 3/2013 | 0/17 | 0/151 hr | 0/17 | |
| Sess 1/2014 | 0/21 | 0/178hr | 0/21 | |
| Sess 2/2014 | 0/28 | 0/275 hr | 0/28 | |
| Sess 3/2014 | 0/16 | 0/160 hr | 1/16 | excluding GP053B-D abandoned |
| Oct-Feb e-VLBI | 0/11 | 0/96.5 hr | 0/11 | |
| Jan'15 OoS | 1/1 | 12/12 hr | 1/1 | GA035B: awaiting Yy pack |
| Sess 1/2015 | 1/17 | 38/272 hr | 16/17 | excluding EG078C abandoned |
| Mar-May e-VLBI | 0/7 | 0/45 hr | 0/7 | |
| Sess 2/2015 | 29 | 330.5 hr | {prognosis} | |
| Jun e-VLBI | 2 | 22 hr | {prognosis} | incl. trigger |

Some landmarks since the previous TOG report:

Session 2/2014
This has proven to be by far the most complicated session we've encountered. There were three large experiments requiring new features on unprecedented scales.  Two of these were distributed in April, one in late May.
The back-log these have established still affects us; below has more details on each, including new operational functionalities that they have engendered.

GP051:  A wide-field spectral line (8192 frequency points over a 16MHz channel, 0.35s integrations) 24-hour observation.  This produced 5.3 TB of output FITS files, and the run-time for post-correlation processing to make/archive the FITS files took 190 hr (correlation itself took only 51hr). The ability to use integration times other than 2^n seconds from the correlation-control GUI was developed for this experiment (going from 0.25s to 0.35s saved about 2.1 TB in output size while maintaining a satisfactory time-smearing field of view).  This experiment was the driving factor behind increasing disk space in the Archive (and the off-site back-up), as well as obtaining a newer-generation LTO drive for the Archive tape back-up.

GP052:  This was our first experiment to need coherent pulsar de-dispersion. The target was a millisecond pulsar (P=1.6ms), with correlation in three passes to achieve three different gating/binnings:  (i) 16 bins over the full period, (ii) 4 bins over the main pulse, and (iii) 1 bin over the

interpulse. The first two used 512 frequency points per 16MHz channel and the third 4096. The FFT size kept the bins in (ii) to ~2% of the period (a bin can't be smaller than the correlation FFT size ). This project required us to install more SFXC memory (the size of the FFT in the coherent de-dispersion corresponds to the dispersion-time across a 16MHz band, or 4 million points in this case), and also led to a few additions to the correlation-control GUI: specification of the number of processes to be run per node (to avoid memory use problems), specification of the spectral weighting function to use (the PI wanted top-hat rather than the default Hanning for the interpulse pass), and the choice of incoherent or coherent de-dispersion, selectable per pulsar target. The total size of the output FITS files was 1.7 TB. Correlation took 137.5 hours.

EG078B: This was our first multiple phase center correlation with more than ~50 targets. The proposal expected to have ~300. There was no action communicating the actual target list until ~mid-December, after a new student became involved in the project. The final list had 699 targets, extending out to 27'. With the current SFXC cluster (after the GP052 memory upgrades), we could provide ~5% time-/BW-smearing at maximum target separation in the "internal" wide-field correlation underpinning the shifting/re-averaging to the individual targets, as opposed to <1% as we have done hitherto. Correlation made use of the new memory-control features in the GUI developed for GP052. With 9 stations, the multiple-phase-center scans correlated at 10.7x real-time, with 10 stations, at 13.3x (scaling reasonably well with the square of the number of stations). The total correlation time was 227.5 hr (a record for SFXC). The total size of output FITS files was 3.8 TB, and the run-time of the post-correlation steps was 361.5 hr. There are two further epochs to come, most likely with 12-13 stations (EG078B observed while Mc and Ur were out for repairs, SRT should be able to join).

The figure at the end shows the growth of FITS-file size for user experiments on the EVN Archive, with the period since the previous TOG highlighted in red. About 40% of the total archive size has resulted from this period, starting with these large experiments from session 2/2014.

With the RDBE back-ends, VLBA stations no longer provide the TSYS section in {exp}cal.vlba files. An 80-Hz continuous-cal system temperature extractor was developed for SFXC. Globals from session 2/2014 were the first to use this new feature. We have modified the EVN pipeline to accept this new input for ANTAB. If more stations could do continuous-cal reliably, then their antabfs files derived from logs wouldn't not be required (though their rxg files would still be).

Session 3/2014

All experiments have been correlated, and all except for GB075B have been distributed. GB075B, a Q-band observation in which Ef had no fringes (nor in the Q-band NME), took longer to correlate because Gb also had stretches of corrupted VDIF frames, and we didn't want to give up on both big telescopes at once. But the Gb data also had synch words in their VDIF headers, which provided the possibility to maintain frame tracking. Without Ef, recovering as much of the Gb data as possible became even more important, so we modified the SFXC VDIF reader to take advantage of the synch words.

GA035A was the first EVN/global + RadioAstron observation to come to JIVE for correlation as per a proposal. It was an L-band spectral-line observation. More precisely, the proposal called for JIVE to correlate ground stations in continuum mode (delay/rate solutions) and in a spectral-zoomed mode with 100 m/s velocity resolution over 0.5 MHz (Zeeman effect analysis), and then for ASC to correlate the large ground stations and RadioAstron in a narrower, higher-resolution mode. This was the first experiment for which we used spectral zooming. GA035B (the K-band partner to the L-band GA035A) ran as an out-of-session observation in late January; we're still waiting on a pack from the JVLA.

Session 1/2015
   Mc and Ur return, Mc as a full-time DBBC station
   Jb and Wb conduct parallel DBBC testing in NMEs (Jb successfully)
   FR022 = fringe test using the new DBBC firmware enabling 2 Gbps recordings
      in a DDC personality with 32 MHz subbands.
        [fringes everywhere, but still not fully analyzed]

Correlation completed for all experiments except GN002A (still awaiting a pack from the JVLA). The next epoch of EG078 (see above) was abandoned by the PI; Ef's L-band receiver had broken without being able to repair in time (observation started on a Sunday evening) and SRT had formatter problems leading to only about 2.5 hours of recording (of 24) at the correct 1 Gbps rate.

Session 2/2015
   A record for the amount of disk-pack capacity in a session: 1660 TB (sum of packs sent by JIVE and Bonn correlators plus amount on-hand at stations to be used). Also a record for the largest single-station/session load: 122.18 TB (Ef and On); previously, 65 TB would have been a big load for a station in a session. (One factor in this session is that there was no 5cm frequency block occupying several days a low bit-rate -- almost all observations were continuum at 1 Gbps).

  Wb conducting parallel DBBC tests
  Ef and On conducting parallel fila10g-FlexBuff tests

USER SUPPORT

  There were thirteen data-reduction visits to JIVE in this period, seven of which were EVN-TNA supported and one of which was by first-time visitor to JIVE and first-time EVN user (the new student involved in EG078).

  JIVE support-group staff continue to contact all PIs once the block schedule is made public, and to check over schedules posted to vlbeer prior to stations down-loading them, helping to prevent avoidable errors in the observations themselves. Now that EVN stations are transitioning to DBBC back-ends on separate time-scales, this scheduling help also provides PIs with experiment-specific template "setini" blocks and station catalogs, which can change from session to session. A new sched release (11.4) came out in April; most session 2/2015 PIs used this (no one used a version earlier than 11.3u1, which is the first time in memory that only two versions of sched were used among the collective PIs). There were 3 first-time PIs in session 1/2015 and eight in session 2/2015 (5 of these 8 are students). In both 2015 sessions, DBBC-stations used a new IF/BBC distribution scheme that enables the new DBBC firmware to provide 2Gbps recordings (test in session 1/2015 described further in the "NETWORK SUPPORT" section; advertised for the June proposal deadline).

  The pipeline was modified to accept ANTAB information for VLBA stations derived from our 80Hz continuous-cal extractor, rather than from the (legacy) {exper}vlba.cal files. ANTAB information for the JVLA has been provided by A. Mioduszewski in Socorro for observations in sessions 2/2014 and 3/2014 (no JVLA in any observation from session 1/2015, several in 2/2015).

NETWORK SUPPORT

  In session 3/2014, there was the first instances of a station observing the wrong version of a schedule since we shifted to the system in which PIs upload their schedules to one location on vlbeer, and stations download checked schedules from another location. It is still not clear how this happened: it was a global with a number of intermediate versions as the JVLA configurations were being iterated with Socorro, but the wrong version used was never in the vlbeer .latest/ subdirectory. The working hypothesis is that it was downloaded from Socorro (since the station was Ef, who needs access there for HSA observations). Later in session 2/2015, another

station warned us of an error in a vex file, but one that had been fixed in the .latest/ subdirectory (thus they must have obtained that one from the PI-upload area). Do stations think that the current procedures for separating PI-upload and station-download using the .latest/ subdirectories need revision?

Session 1/2015 had a 2 Gbps test using the new DBBC firmware version v105E. This provides BBC filters ranging from 2 to 32 MHz (as opposed to the current firmware giving 1 to 16 MHz), but can use only the odd-numbered BBCs. Four core2 boards are thus needed to get the eight dual-sideband BBCs needed to attain 2 Gbps; some stations have only two core2 boards, so even with the new firmware would be limited to 1 Gbps. The new firmware also leads to a change in the IF patching. Previously, four dual-sideband BBCs in each polarization came from one conditioning module (RCP=A, LCP=B); with the v105E firmware this needs to change to RCP=A+B, LCP=C+D (ignoring details related to "astro"/"hybrid"/"geo" configurations). To avoid having to repatch the LCP channels between experiments, all observations in sessions 1-2/2015 were scheduled using patching consistent with the v105E case. Initial results from the 2 Gbps test (FR022) shows success at all stations, but with phase across the bands less linear than in the Gbps case in the current firmware. Detailed investigation is continuing.

In sched, at the request of Onsala we shifted our session-specific stations catalog and corresponding setini plug-ins to use abbreviations "O8" and "O6" for station-names ONSALA85 and ONSALA60 (i.e., using feet as the diameter unit), while retaining the abbreviation "On" for both "new" station-names ONSALA25 and ONSALA20 (i.e., using meters). This facilitates using both antennas in the same observation as distinct stations, while retaining the ability to use the traditional 'On' abbreviation. Craig added explicit pointing-sector in late January 2014; currently this is used for only VLBA stations. A separate e-mail last week discusses possibilities for EVN stations to get similar control. Now that almost all stations have DBBC-compatible patchings, the time seems ripe to enter these into the normal sched catalogs. I still have to work out the issue of the v105E complications (mentioned above) for stations currently without four core2 boards and for the Chinese stations, both of whom have not had to worry about this yet. In order to give full flexibility for DBBC firmwares both with and without the 32MHz filter capabilities, we will likely need to introduce a new control parameter, since there will be different patching rules (e.g., old = 1-16MHz filters, all BBCs; new = 2-32MHz filters, only odd BBCs). This is finer than a pure station-based parameter that can live in the stations catalog (as is DBBCVER = astro/geo/hybrid), so the most efficient and/or transparent parameterization may take more reflection.

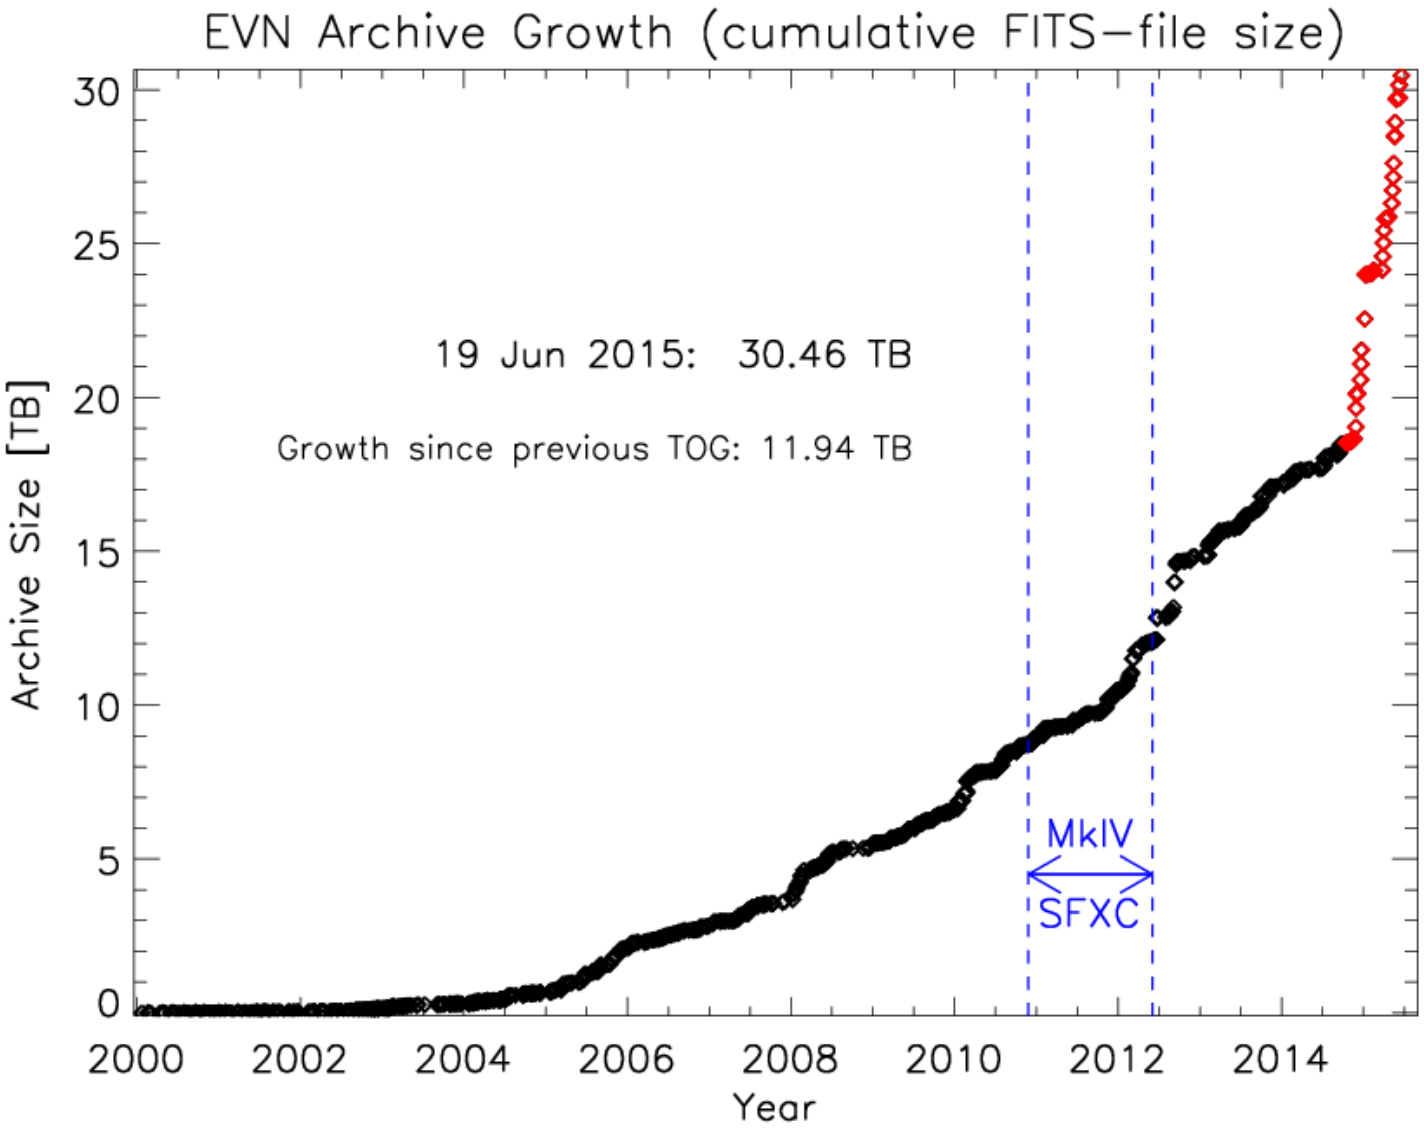
EVN Archive Growth (cumulative FITS−file size)
19 Jun 2015: 30.46 TB
Growth since previous TOG: 11.94 TB
MkIV
SFXC
Archive Size [TB]
Year
0
5
10
15
20
25
30
2000
2002
2004
2006
2008
2010
2012
2014