PROCEEDINGS OF SCIENCE

e-VLBI and Other Developments at the EVN Data Processors at JIVE

R.M. Campbell[*†]
Joint Institute for VLBI in Europe
E-mail: campbell@jive.nl

We review the capabilities of the EVN Mk IV Data Processor at JIVE, with special emphasis on the new recirculation feature that provides higher spectral resolution for observations using ≤ 8 MHz subbands that would otherwise have been limited by having the number of baselines or polarizations. We then move to the real-time e-EVN facility that has matured over the past couple years to the point where sustained 8–9 station Gbps observations are routine (from the PI's perspective...). With the advent of such a real-time array, the number of target-of-opportunity observations has proliferated, clearly having touched a pent-up demand. We conclude with a glimpse into the new SFXC software correlator at JIVE, which has been processing ftp fringe-test observations within Network Monitoring Experiments for more than three years, and recently completed correlation of its first user experiment that required pulsar gating.

The Science Operations & Support Group at JIVE continues to provide assistance for all phases of user experiments — from proposing and scheduling, through correlation, to analysis of the resulting FITS data. Campbell (2008 [1]) presents a more complete review of the resources available and the pre- and post-correlation operational flows. The EVN web-page and the (revised) JIVE web-page are the broadest sources of information. The EVN Users' Guide provides the best "first-stop" for on-line help, with links to more explicit help geared towards the specific tasks encountered in conducting EVN experiments (proposing, scheduling, correlating, analyzing). Almost all EVN and global proposals are now submitted through the EVN proposal tool (target-of-opportunity and short observations are the exceptions). The EVN Archive contains FITS files, standard plots, station feedback, and pipeline results, with security provided in accordance with the EVN data access policy. In order to facilitate visits to JIVE or other EVN institutes, the EC provides a trans-national access program for eligible experiments. See the Access to the EVN web-page for the latest details.

10th European VLBI Network Symposium and EVN Users Meeting: VLBI and the new generation of radio arrays
September 20-24, 2010
Manchester, UK

[*]Speaker.
[†]on behalf of the EVN Data Processor Group at JIVE.

 http://pos.sissa.it/

PoS(10th EVN Symposium)034

1. Capabilities of the EVN Mk IV Data Processor

The EVN Mk IV data processor can correlate simultaneously up to 16 stations with 16 channels per station, *i.e.*, either 8 dual-pol or 16 single-pol subbands (SBs), each having a maximum sampling rate of 32 Msample/s (16 MHz Nyquist-sampled $BW_{\rm sb}$). This provides for a total recording rate of 1 Gbps per station using 2-bit sampling (VLBA antennas can currently participate in global observations with EVN stations recording 1 Gbps by using 1-bit sampling at 512 Mbps, to maintain uniform $N_{\rm sb}$ and $BW_{\rm sb}$ throughout the array). We can currently correlate/provide: (*i*) Mark 5A or Mark 5B recordings (5B playback via 5B or 5A+); (*ii*) up to 2048 frequency points per baseline/subband/polarization, *cf.* §1.1; (*iii*) oversampling at 2 or 4 times the Nyquist frequency, in order to provide subband bandwidths down to 500 kHz; (*iv*) full-correlator integration times down to 0.25 s, *cf.* §1.2; and (*v*) real-time e-VLBI operation, *cf.* §2.

1.1 EVN Mk IV Correlator Capacity

The single-pass Mk IV correlator capacity can be expressed:

$$N_{\rm sta}^2 \cdot N_{\rm sb} \cdot N_{\rm pol} \cdot N_{\rm frq} \leq 131072 \cdot \mathscr{R}. \qquad (1.1)$$

Here, $N_{\rm frq}$ is the number of frequency points per baseline/subband/polarization. $N_{\rm pol}$ is the number of polarizations in the correlation (1, 2, or 4), and $N_{\rm sb}$ is the number of subbands, counting lower- and upper-sidebands from the same BBC as distinct SBs. The value to use for $N_{\rm sta}$ is “granular" in multiples of 4 (*e.g.*, if you have 5–8 stations, use 8). Independent of equation 1.1, the maximum number of input channels ($N_{\rm sb} \cdot N_{{\rm pol}_{\parallel}}$) is 16, and the maximum $N_{\rm frq}$ is 2048 (a single interferometer must fit onto a single correlator board).

Recirculation time-shares correlator chips in the case that the observations don’t require the correlator to operate at its maximum rate (*i.e.*, $<$ 32 Msamples/s, or $BW_{\rm sb} < 16$ MHz). This effectively increases the correlator capacity for such experiments: more lags can be processed in what otherwise would have been idle correlator-chip cycles. The recirculation factor in equation 1.1 is $\mathscr{R} = 16\,{\rm MHz}/BW_{\rm sb}$ (for Nyquist sampling), up to $\mathscr{R}_{\rm max} = 8$. However, the maximum $N_{\rm frq}$ would remain 2048, as discussed above. The principal beneficiaries of recirculation would be line experiments whose spectral resolution would be limited by $N_{\rm sta}$ or $N_{\rm pol}$ — then EVN now has reliably >8 stations able to observe OH or methanol maser emission, and interest in full-Stokes mapping of these masers is growing. Some experiments may prefer 16 MHz subbands at lower data rates (*e.g.*, extra-galactic HI absorption); these would not be able to gain from recirculation.

The combination of $N_{\rm frq}$ and $BW_{\rm sb}$ sets the velocity spacing of an observation/correlation in terms of velocity per frequency-point. The fundamental equation is:

$$\Delta {\rm v} \ = \ c \, \frac{BW_{\rm sb}/N_{\rm frq}}{\nu_{\rm obs}}\,. \qquad (1.2)$$

Table 1 shows some configurations that would require the full correlator capacity to configure. Note that with recirculation, the correlator capacity has become a function of $BW_{\rm sb}$. Table 2 shows velocity spacings provided by the maximum $N_{\rm frq} = 2048$ for various combinations of spectral lines and $BW_{\rm sb}$. Recirculation and oversampling have been shown not to work well together, thus

N_{sta}	N_{sb}	BW_{sb}	N_{pol}	N_{frq}	$\mathscr{R}$	comment
5–8	1	16	1	2048	1	EVN spectral-line (no Recirculation)
5–8	1	16	4	512	1	with cross-pol
9–16	1	16	1	512	1	global spectral-line (no Recirculation)
9–16	1	8	2	512	2	new modes with Recirculation
9–16	1	2	2	2048	8	
9–16	1	2	4	1024	8	
9–16	8	16	4	16	1	global continuum Gbps cross-pol

Table 1: Examples of "maximal" correlator configurations.

BW_{sb} [MHz]	$H_{(1420)}$	$OH_{(1665)}$	$CH_3OH_{(6668)}$	$H_2O_{(22235)}$	comment
16	1651	1408	351	105	
2	206	176	44	13	
0.5	*52*	*44*	*11*	*3.3*	*no Recirculation*

Table 2: Velocity spacing [m/s] for $N_{frq} = 2048$, for various spectral lines and BW_{sb}.

$BW_{sb} < 2$ MHz could not use recirculation. Since $\mathscr{R} = 8$ for 2 MHz subbands, recirculation provides a greater boost to total spectral capacity than does oversampling — but cannot achieve velocity spacings below those of the 2-MHz row in table 2 because of the hard maximum of $N_{frq} = 2048$ per baseline/SB/polarization.

1.2 Output Capacity

The minimum t_{int} for a configuration using the whole correlator is $1/4$ s; some modes using no more than half the correlator can achieve $1/8$ s. With recirculation, the minimum t_{int} would be increased by a factor of $\mathscr{R}$ from its nominal value. The recirculation choices at correlation are $\mathscr{R} = 16\text{MHz}/BW_{sb}$ or $\mathscr{R} = 1$ (*i.e.*, recirculation off) — thus requiring short t_{int} may sacrifice some spectral resolution for narrow-bandwidth experiments. FITS files resulting from maximum correlator-output rate grow at about 7–10 GB per correlator pass per hour of observation. The record for the largest total size of output FITS files for a single epoch of an experiment remains 1028.7 GB, but data-sets in the range of a few hundred GB are growing more common.

The combination of short t_{int} and large N_{frq}, provide a wide-field mapping capability through reduced bandwidth- and time-smearing effects in the *u*-*v* plane. The EVN calculator evaluates field-of-view limitations using the formulation of Wrobel (1995 [9], §21.7.5). There is also a perhaps non-intuitive conflict between sensitive 1 Gbps recordings and the field-of-view limit from bandwidth-smearing. A Gbps rate implies that $N_{sb} \cdot N_{pol}$ in equation (1.1) will be 16 (no cross-pols) or 32 (with cross-pols), and BW_{sb} will be 16 MHz (thus no gain from recirculation). An 8-station Gbps experiment without cross-pols can get at most $N_{frq} = 128$ in a single correlator pass, with a resulting $FoV_{BW} \sim 6\overset{\prime}{.}6/B_{1000}$, where B_{1000} is the longest baseline length in units of 1000 km. This would be (cumulatively) halved if using cross-pols and quartered if $N_{sta} \geq 9$.

PoS(10th EVN Symposium)034

2. Real-time e-VLBI in the EVN

Real-time e-VLBI has become a proposal-driven operational mode within the EVN, with support from the EXPReS (2006–9) and NEXPReS (2010–present) projects funded by the EC (DG-INFSO). The first proposal-driven e-EVN observation was on 16 March 2006, using six stations at 256 Mbps. Szomoru (2008 [8]) and Campbell & Szomoru (2009 [2]) provide a review of e-EVN developments through September 2008 and September 2009, respectively. The combined effects of technical advances in the networks, stations, and correlator have enabled routine sustained 1024 Mbps real-time observations for some time now. Ef, Wb, Jb, On, Tr, and Mh can transmit a full 1024 Mbps data rate. However, this can not fit through a standard 1 Gbps fibre link; stations currently in such a situation (Mc, Ys, Hh) can participate via “channel-dropping" operating in the station’s Mark5 unit — in this case, 7 of the 8 dual-pol subbands are transmitted without loss and the remaining one is not sent, providing a solution to reducing the bandwidth to fit into the available fibre that is more predictable than was the earlier packet-dropping. Transmission from Jodrell Bank has included 1024 Mbps from a home station plus up to four (microwave-linked) MERLIN out-stations at 128 Mbps each; in principle we should be able to achieve 1024 Mbps and 768 Mbps from the two home stations (*i.e.,* full Jb1 plus 6/8 subbands from Jb2 for help in phase-steering, with channel-dropping from Jb2). Arecibo and Sheshan are limited to 256 Mbps for now, with Ar able to reach 512 Mbps during early morning (local) hours. The e-VLBI status table provides the current capabilities and scheduled observing dates of the e-EVN array.

The gains in the number of e-capable stations and the achievement of Gbps data rates have provided a highly competitive e-EVN array in terms of sensitivity and u-v coverage. The key distinctions of the real-time e-EVN are the extremely reduced latency in receiving correlation results — we can provide final FITS files within a day of the end of observations, in some cases the same day — and the more frequent observing opportunities. For sources that vary on short time-scales (flaring X-ray binaries, gamma-ray bursts and other transients, just-exploded supernovae) such short-latency high-resolution VLBI results can be vital to adapting observing tactics based on the source’s behavior. Increased density of observing epochs beyond the usual thrice-annual EVN sessions also provides the potential for better synchronization of VLBI observations with campaigns at other wavelength bands, or for better monitoring of kinematics or population changes.

e-EVN observing sessions are currently scheduled for 24 hours on pre-arranged dates, about once per month. There have been a few paired dates 2–3 days apart to permit following source changes on shorter time-scales. Proposals for e-EVN observations are now submitted in the same fashion as are other EVN/global proposals at the thrice-annual deadlines, and compete for the available time with their peers based on the grades of the EVN programme committee. A class of “triggered" proposal enables e-EVN observations of a pre-selected source when its behavior prior to an e-EVN session shows that it has entered an interesting state. Once such a triggered proposal has been accepted, the proposing group only needs to submit a short trigger request up to 24 hours before the start of an e-EVN session to be considered for observation in that session. There have been six such triggered observations. As always, target-of-opportunity (ToO) proposals may be submitted at any time following the usual EVN procedures. The development of the high-sensitivity real-time e-EVN network has spurred interest in such ToOs, to judge from their increasing numbers.

Since the beginning of proposal-driven e-EVN observations through the end of November

PoS(10th EVN Symposium)034

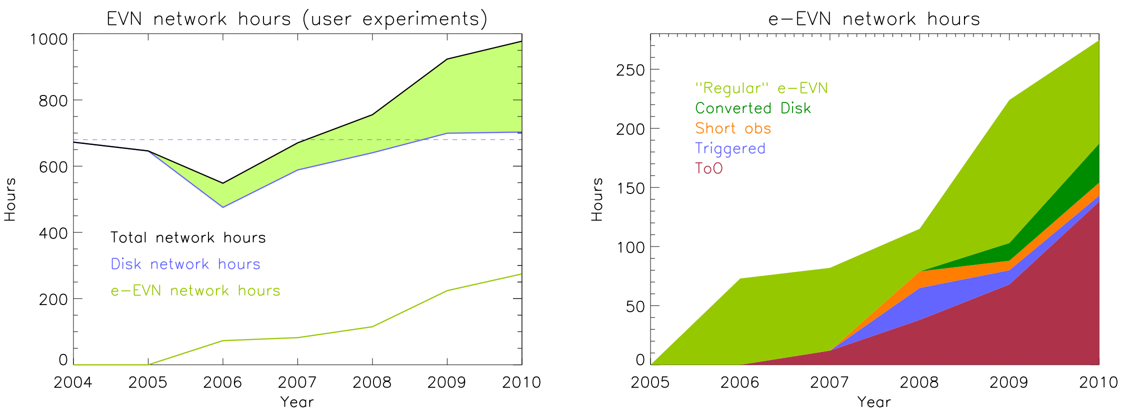

Figure 1: Growth of e-EVN observing. Left-hand panel: annual EVN network hours, with e-EVN contribution in green. Right-hand panel: evolution of annual network hours by type of e-EVN observation.

2010, there have been 97 observations from 67 proposals for a total of 768.5 hours. These observations represent 32 different PIs. The left-hand panel of figure 1 shows the growth of e-EVN observations on top of the traditional tape/disk-based observations (annual binning). Since 2004, the number of recorded EVN network hours has not deviated much from about 680 hours (the horizontal dashed line), barring a reduction in 2006 stemming from maintenance work on Effelsberg. e-EVN observations have grown to 274.5 hours so far in 2010 (with one further scheduled e-EVN day to come in 2010), more than 28% of the total observing time. The right-hand panel of figure 1 shows the evolution of the composition of the e-EVN observations and the growing importance of real-time ToO experiments – so far in 2010, more than half the total e-EVN observing time (138 hours in 15 observations from 11 proposals). Some ToOs have also included Australian and Japanese stations in the real-time correlation (at 512 Mbps).

3. Software Correlation on SFXC

SFXC is an FX-based distributed software correlator based upon the correlation algorithms developed for observing the descent of the Huygens probe onto Titan (Pogrebenko et al. [6], [7]), adapted to wide-field applications. Its development is/has been partly covered under EXPReS, NEXPReS, and the SCARIe projects. It currently runs on a dedicated 16-node, 128-core cluster, with quad-rate Inifiband (40 Gbps) connections among the nodes. This suffices to correlate a 512 Mbps experiment on a 9-station array in real time. We hope to be able to at least double the size of the cluster in early 2011. See Kettenis et al. (2010 [4]) for more technical details.

From an astronomical viewpoint, SFXC enables observations that are not currently possible on the EVN Mk IV correlator. Principal among these are pulsar observations with gating and binning. SFXC can provide an arbitrary number of equally-spaced bins within a specified gate (fraction of a pulse-period), each bin resulting in a separate output correlation product. Traditional pulsar gating as a means to boost SNR on pulsar detections (to use, say, for astrometry) would correspond to one bin. We have completed correlation of the first user experiment on SFXC including the pulsar gating capability in November 2010. SFXC also avoids the hard spectral- and output-rate limitations discussed in §1.1 & 1.2, thus has a natural affinity for spectral-line experiments needing

PoS(10th EVN Symposium)034

better velocity spacings than in tables 1 and 2, and for wide-field mapping experiments. As an FX correlator with true (analytic) station-based fringe-rotation, SFXC also provides better phase-noise characteristics; Pidopryhora et al. (2009 [5]) present some comparisons of the Mk IV, SFXC, and DiFX correlators. A part of NEXPReS addresses integration of SFXC with the real-time e-EVN observations, and also the possibilities for globally distributed correlation (Kettenis et al. 2009 [3]).

We have been using SFXC (on a smaller cluster) since June 2007 to process the ftp fringe tests that take place during the Network Monitoring Experiments in each frequency sub-session of EVN sessions. An automatic-ftp feature in the Field System copies a specified portion of a scan from the station's Mark5 pack to a linux file and ftp's it directly to JIVE, where the arrival of new data is detected, correlation performed, and results posted to a web page available to the stations. These ftp fringe tests have been very successful in identifying problems early enough to allow stations to repair them before user experiments would have been affected.

Acknowledgments

The European VLBI Network is a joint facility of European, Chinese, South African and other radio astronomy institutes funded by their national research councils. This effort/activity is supported by the European Community (EC) Framework Programme (FP) 7, Advanced Radio Astronomy in Europe, grant agreement No. 227290. EXPReS was an integrated infrastructure initiative (I3) funded under the EC FP 6, grant agreement No. 026642. NEXPReS is an I3 funded under the EC FP 7, grant agreement No. RI-261525. SCARIe has been funded under subsidy No. 643.200.504 from the Nederlandse Organisatie voor Wetenschappelijk Onderzoek (NWO).

References

[1] R.M. Campbell, *Recent Results from the EVN Mk IV Data Processor at JIVE*, in proceedings of *The 9th EVN Symposium*, 2008, `PoS(IX%20EVN%20Symposium)042`.

[2] R.M. Campbell & A. Szomoru, *e-VLBI and Other Developments at the EVN MkIV Data Processor at JIVE*, in proceedings of *The 19th European VLBI for Geodesy and Astrometry Working Meeting*, 79, 2009, `www.u-bordeaux1.fr/vlbi2009/proceedgs/19_Campbell.pdf`

[3] M. Kettenis, A. Keimpema, D. Small, & D. Marchal, *e-VLBI with the SFXC Correlator*, in proceedings of *The 8th International e-VLBI Workshop*, 2009, `PoS(EXPReS09)045`.

[4] M. Kettenis, *SFXC: a Distributed Software Correlator for VLBI*, in proceedings of *The 10th EVN Symposium*, 2010, `PoS(10th%20EVN%20Symposium)086`.

[5] Y. Pidopryhora, A. Keimpema, & M. Kettenis, *The Latest Tests of the SFXC Software Correlator*, in proceedings of *The 8th International e-VLBI Workshop*, 2009, `PoS(EXPReS09)046`.

[6] S.V. Pogrebenko et al., *VLBI Tracking of the Huygens Probe in the Atmosphere of Titan*, `www.mrc.uidaho.edu/entryws/full/programme_detailed.html`, C-4.6

[7] S.V. Pogrebenko et al., *JIVE Research Notes #4, 5, 11*, `www.jive.nl/jive-research-notes`.

[8] A. Szomoru, *EXPReS and the e-EVN*, in proceedings of *The 9th EVN Symposium*, 2008, `PoS(IX%20EVN%20Symposium)040`.

[9] J.M. Wrobel, *VLBI Observing Strategies*, in *VLBI and the VLBA*, eds. J.A. Zensus, P.J. Diamond, & P.J. Napier, ASP, San Francisco, 411, 1995.