JUMPING JIVE

730884 – JUMPING JIVE – H2020-INFRADEV-2016-2017/H2020-INFRADEV-2016-1

Software to deal with geodetic observing schedules

Deliverable 6.2

Submission date: 29.11.2017

I. Background

A. Astronomical VLBI

The European VLBI network (EVN) conducts astronomical VLBI observations, for the vast majority of which the experimental design comprises the full array tracking a specific target across the sky for a period of time. For faint targets or experiments that aim to measure differential astrometry, this tracking interleaves observations of the target and a phase-reference calibrator that lies nearby on the plane of the sky on time-scales ≤10 minutes, in order to calibrate out effects of propagation through the atmosphere from the observations of the target. All network stations would have schedules in which they all observe the same set of sources in a synchronized fashion via a series of "scans" (defined here as a block within a schedule having a specific set of stations, source, and time-range). In this fashion, the image sensitivity generally improves the fastest, as the collecting area of all network antennas contributes at once.

Very occasionally, scheduling efficiency can be increased by splitting the full array into different sub-arrays, with each sub-array tracking a different target (or target + phase-reference source) over a time-range. This has happened only a few times, typically in experiments having multiple targets at low declinations well-spread in right ascension and an array with large east-west extent. The combination a network with large east-west extent and low-declination sources implies that not all antennas can see the targets at the same time (i.e., the target would be below the horizon for easternmost or westernmost part of the overall network). For example, in a global VLBI observation (EVN + the VLBA in the U.S.), different sub-array configurations could be {EVN} and {VLBA} or {EVN + eastern-VLBA} and {western-VLBA}. In cases like this, there would be a time-range in which each sub-array would observe independently of the other, but within each sub-array the concept of synchronized observing as described in the first paragraph would still hold.

B. Geodetic VLBI

There is a principal distinction between the experimental design for astronomical VLBI described above and for VLBI used for geodesy or absolute astrometry as organized by the International VLBI Service for Geodesy and Astrometry (IVS). In geodetic VLBI, the (extra-galactic) sources serve as beacons in an inertial frame from which to determine the position of the observing antennas. There is no phase-referencing with which to calibrate out the effects of propagation through the atmosphere, so instead the schedule prioritizes observing as many sources spread over the whole visible sky as quickly as possible, to be able to estimate parameters related to atmospheric propagation together with antenna positions in a single fitting process. The concept of a "scan" here is much more local than for standard astronomical VLBI — maximizing the number

of sources spread throughout the sky leads to a much more fluid situation of constantly changing sub-arrays per scan, scan boundaries overlapping in time, and individual stations leaving a scan prematurely to join a different scan with other antennas, perhaps even before its original scan has finished. A VLBI schedule for absolute astrometry (i.e., determining positions of sources within a celestial reference frame) will share many of the characteristics of a geodetic-VLBI schedule, because the central aim of estimating atmospheric-propagation parameters from the VLBI observations themselves remains in place.

II. The EVN software correlator at JIVE (SFXC)

Because the EVN software correlator at JIVE (SFXC) was designed to process EVN astronomical VLBI observations, one of the underlying assumptions in its control system was that time-range (i.e., start time, end time) could be used as a unique identifier for the composition of a correlation job[1]. SFXC is driven by information in two files: the observing schedule (from which it also obtains, among other things, the coordinates for the antennas and sources and information about how each antenna's data are encoded) and a "*ctrl*" file, which passes along specifics of an individual correlation job —including the start and end time and the list of antennas plus the location of their data). Given these two defining files, SFXC will then compute the *a priori* geometric delay model, perform the correlation, and output the complex visibilities residual to the *a priori* model. In standard operations, the interface program "*runjob*" sits between the correlator operator and SFXC itself. This program presents information from the observing schedule in a GUI display, allowing the operator to select the time-range, stations, and channels to include in the job, as well as to specify the correlation parameters, such as the frequency and time resolution to use. Once the job is started, the *runjob* program makes the *ctrl* file discussed above.

In cases of sub-netting as discussed in Section I, some of the defining properties of a job were left ambiguous by using only the start-/end-time as the fundamental criterion. Depending on the precise form of the sub-netting in question, effects could range from not using the appropriate source for computing the *a priori* model to the correlator hanging when an overlap period between scans ends. Since such sub-netted schedules have occurred so infrequently in EVN observations, up until now we have handled such situations by making two copies of the actual observing schedule, removing the minimum number of scans from one copy to avoid having more than one scan running at any given time, and retaining only those scans in the other copy. This of course requires running both schedules as separate entities, and later re-combining the output data in the post-correlation analysis stage.

[1] Here, we ignore the operational distinction between correlator job and sub-job.

To overcome these limitations, we adapted both the SFXC code and the *runjob* program. For SFXC, a optional scan label was added to the low-level selection criteria. If passed, it would break the ambiguity among multiple scans in the observing schedule that overlap in time. SFXC can further establish antenna-specific start/end time-ranges that correspond to the specified scan, and can flag any data from an antenna that lies within the "global" start/end time-range but outside of its own one. This would preclude spurious data entering into the correlation output. In order to pass this finer level of job definition along to SFXC, *runjob* now outputs the scan-label from the observing schedule to the *ctrl* file, and also does not raise objections when it encounters sub-netting within an observing schedule.

This ability to handle a sub-netted schedule transparently within a single correlator job exists within release 3.5 of SFXC. The SFXC page on the JIVE wiki:
http://www.jive.eu/jivewiki/doku.php?id=sfxc
contains a link to the SVN code repository:
https://svn.astron.nl/sfxc/branches/stable-3.5/

III. Testing

We tested the operation of these new SFXC/*runjob* capabilities using data from the experiment R1680, one of the series of 24-hr geodetic experiments conducted each Monday (the "1" in the experiment name) by the International VLBI Service. It observed on 23 March 2015, and was correlated through the usual IVS programme. Data from four stations that were still available at the time served as the basis for the successful "preliminary assessment of the astrometric quality of SFXC" that was mentioned in the WP.6 description in Table 3.1.a of the Jumping JIVE proposal. For that assessment, we processed the sub-netting in the manner described in Section II: creating two non-sub-netted schedules out of the original sub-netted one, and correlating both in turn.

For the purpose of testing, we identified the hour with the highest SNR targets from the original processing whose scans were involved in sub-netting (among the 501 total scans involving two or more of the four stations, 134 were involved in sub-netting situations). We then correlated them again in the SFXC stable 3.5 release from the original schedule in one job. Figure 1 below lists the four antennas and illustrates the pattern of scans over this time-range, both in the original schedule (upper half) and in the split schedules that were correlated previously (lower half). Here, the instances of sub-netting are circled in the plot of the sub-netted schedule. The narrow lines in the plot of the split-apart schedules denote scans moved from the original schedule into the second one in order to enable correlate before the SFXC/*runjob* changes described here. Further, comparison of the pattern of scans from the original and the split-apart

schedules illustrates another adjustment required to enable correlation before SFXC release 3.5: due to the use of only four of the ten antennas from R1680, some scans had one antenna having a longer duration within a scan than any other antenna —the split-apart schedules shortened these to the duration of the next-to-longest antenna (cf. Ma in the scan starting just after 08:24, Wz in the second sub-netted event, Ny in the third, or Ma in the last). Such an additional adjustment to the schedule passed to the correlator was not required with the new SFXC version.

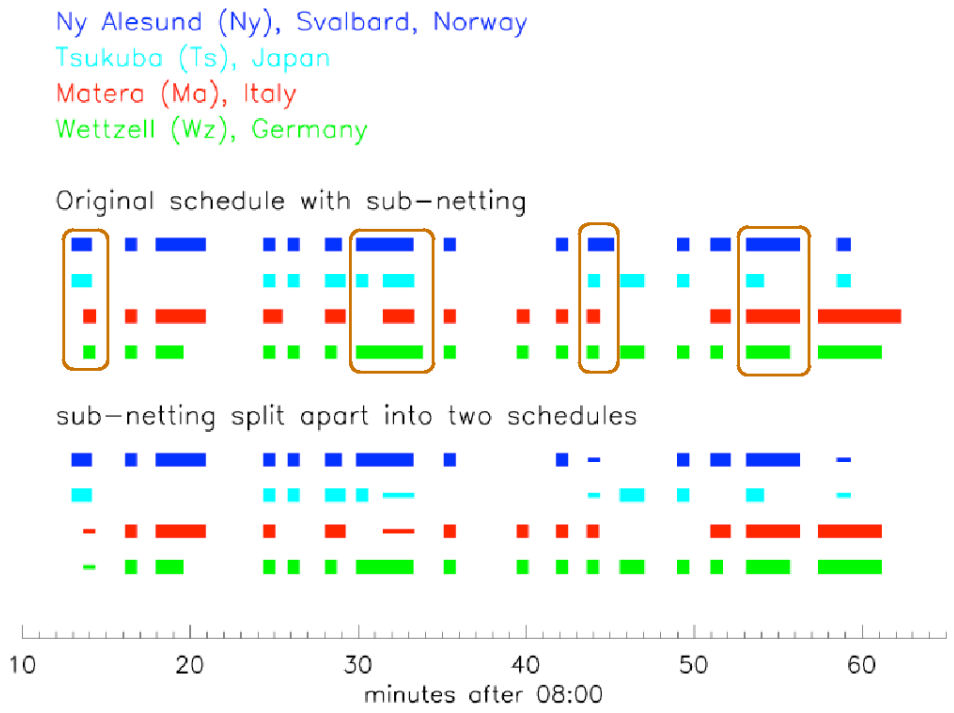

Figure 1. Schematic of scans in the geodetic-VLBI experiment R1680 over the time-range used for comparisons shown in this report. Scans with sub-netting among the four antennas in the original schedule are circled in brown.

In the second instance of sub-netting, one station (Ts) leaves the first scan, and begins a new scan with a different station while the first scan is still underway. Figure 2 shows these two scans extracted from the four-station schedule. The first scan here begins at 08:29:53 on the target 0955+476, with Ny and Wz tracking this source for 204 and 239 seconds respectively (i.e., the Wz behaviour mentioned in the previous paragraph). However, Ts remains on this target for only 43 seconds, and then joins Ma in a different scan at 08:33:26 observing a different target (1842+681).

```
  scan 083-0829;   start = 2015y083d08h29m53s;  source = 0955+476;
    station = Ny :    0 sec :   204 sec :     0 GB :    :    : 1;
    station = Ts :    0 sec :    43 sec :     0 GB :    :    : 1;
    station = Wz :    0 sec :   239 sec :     0 GB :    :    : 1;
  endscan;
*
  scan 083-0831;   start = 2015y083d08h31m26s;  source = 1842+681;
    station = Ma :    0 sec :   114 sec :     0 GB :    :    : 1;
    station = Ts :    0 sec :   114 sec :     0 GB :    :    : 1;
  endscan;
```

Figure 2. Extract from the original four-station R1680 schedule for the second instance of sub-netting (cf. Figure 1).

Correlation in SFXC stable release 3.5 from the original sub-netted schedule did proceed successfully (as an extra check, a similar attempt with release 3.4 hung while processing the second scan of the first instance of sub-netting, at the end-time of the first scan —note that the duration of the Ma-Wz scan extends 15 seconds beyond the end of the Ny-Ts scan, as can be seen in Figure 1). Figure 3 shows the residual phase(*t*) plot resulting from the single-job correlation through release 3.5 (left-hand panel) and similar plots from the earlier two-job correlation (right-hand panels).

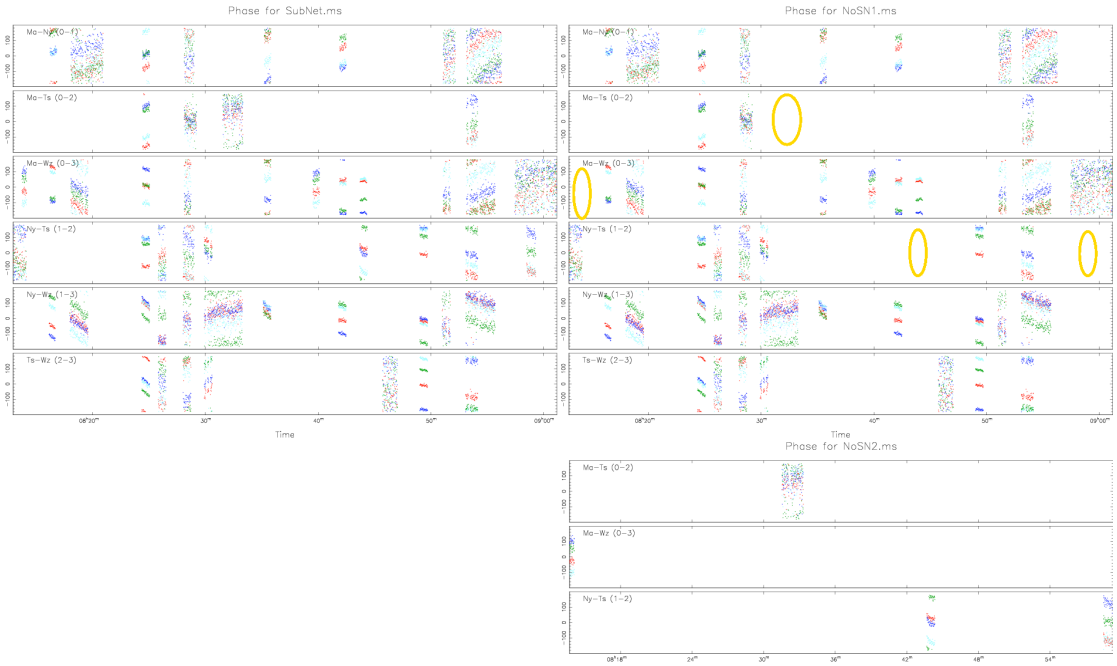

Figure 3. Phase(*t*) of output visibilities (only 4 channels plotted) for the cases of the new correlation using SFXC release 3.5 in a single job (left-hand panel) and the previous correlation run in two jobs after splitting apart the sub-netted scans into different schedules (right-hand panels).

For clarity of the plots, this figure shows only four of the 16 observed channels; each channel is color-coded the same in all plots. The scans split apart into a separate schedule in the previous correlation are marked with yellow ellipses in the top-right panel. In the previous correlation, these

were picked up in the second pass of the split-apart schedule (bottom-right plot), but with the new SFXC release, they are correlated transparently within a single correlator job of the full schedule. Comparing the phase(*t*) values shows that the changes in correlator-control to enable processing sub-netted schedules has not affected the correlator output.

Figure 4 shows similar plots for the weight and phase as a function of time for the scans involved in the second instance of sub-netting, here plotting only one channel, providing a greater resolution into the correlator output. These plots use the same left-/right-panel lay-out as in Figure 3 for the new correlation using SFXC release 3.5 and the previous correlation that had to split apart the sub-netted scans into separate schedules and correlator jobs.

The weight(*t*) panels show that the baselines to Ts in the first of the two scans are not affected when Ts returns with Ma in the next scan that begins before the first scan has finished (the weights stay 0, as data from Ts subsequent to its departure from the first scan are flagged for the duration of that scan). Otherwise, the same pattern of data drop-outs lasting a few seconds at the various antennas remains the same for both the new and old correlation (as these are characteristics in the data themselves). The corresponding phase(*t*) panels permit a more detailed integration-by-integration comparison of the output residual phases, confirming that there are no differences brought about by the changes to the correlator-control logic enabling sub-netted schedules to be processed transparently.

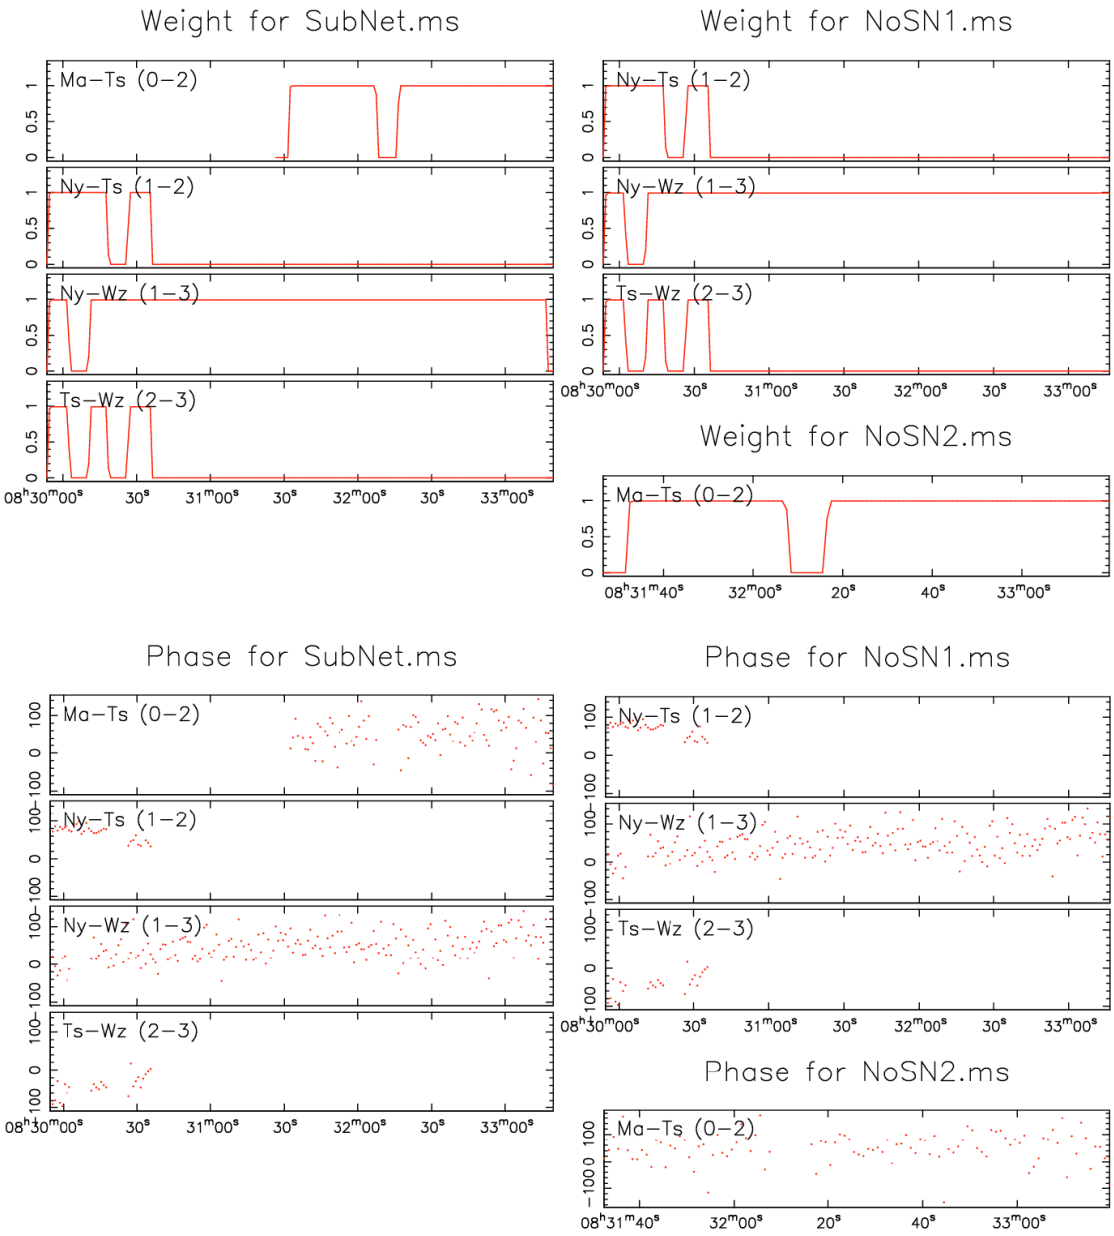

Figure 4. Weight(*t*) and phase(*t*) of output visibilities for the second instance of sub-netting (only one channel plotted). The left-/right-hand panel lay-out matches that of Figure 3.