Dynamic Observing in the VGOS Era

Jim Lovell
University of Tasmania
Draft no. 3, 9 Jan 2014

Introduction

One of the aims of VGOS is for continuous VLBI observations. This is a significant departure from present IVS sessions which are typically 24 h in length for EOP and astrometric observations at the rate of about 170 sessions per year and one hour for UT1 measurements every day of the year. Present sessions are scheduled in advance for an entire year and fix in antenna, media, correlator, scheduling and analysis resources.

One of the goals of VGOS is for initial data products to be made available within 24 h of observation. It is hoped then that communication network capacities will have increased to allow stations to transfer data to the correlators in real time, doing away with the need for media and thus significantly reducing operating costs. Real time correlation also permits immediate feedback to the stations on their performance, allowing them to address any issues (hardware, software, pointing, RFI etc) quickly. This feedback also has the potential to allow for dynamic scheduling where, for example, sources can be rejected after their first observation if they prove unsuitable due to weaker than expected flux density. Observing bands can be shifted to avoid RFI. If an antenna is dropped due to poor performance or wind stow, the schedule can be re-optimised on the fly. Furthermore, data from the correlator could be analysed in real time so that a scan could be observed for just the right amount of time to reach a target SNR.

Another measure to decrease operations costs has been previously discussed where there would be several VGOS Operations Centres (OCs) distributed globally. Staff at an OC would initiate an observation during their business hours and then hand over to another OC to the west at the end of their day. Sharing and transfer of antenna control is already possible, and has been successfully tested and used with the eRemoteCtrl software (Neidhardt et al., 2010).

Under VGOS, there may not need to be a block schedule created for an entire year. Instead, network stations could make their antennas available whenever possible, and the OC could add them to the current session depending on need. The OC has in effect a pool of antennas to draw from at any given time and can select an array appropriate to the desired observing program. Antennas can be added and removed from a session dynamically. A slightly more structured arrangement may be required however to ensure the minimum requirement for an observing program. For example, for intensive UT1 observations, a minimum commitment of two antennas with a long east-west baseline may need to be made ahead of time but the OC could choose to add more antennas if they would improve the solution.

Like the antennas, a correlator could also be treated as an allocatable resource. For example, a schedule aimed at determining a full EOP solution will frequently divide the available antennas into two or more sub arrays. Data from these subarrays could be distributed to different correlators to share the network load. No baselines would be lost as there would be no baselines in common.

The advantages of this type of dynamic observing scenario is that only the required on-source time to reach a target SNR is used, thus achieving the maximum possible number of useful scans per day. Only suitable sources are kept in the observing list and unsuitable ones rejected. Also, and more importantly, antenna-related problems are identified quickly and schedules can be re-optimised. Real-time correlation also potentially allows for real time analysis, reducing the latency for UT1 and full EOP measurements.

In the following section I describe a possible VGOS observing scenario, then in Section 2 I discuss some of the technology and software developments that are still required to achieve this. In Section 3 I present some suggestions on how a fully dynamic observing system might be developed over the next few years as we transition from legacy systems to VGOS and from low to high network bandwidths.

1. The dynamic observing cycle

A dynamic VGOS observing session would be divided into three components: setup, observing and tear-down. It should be possible for more than one observing session to be carried out simultaneously. It is envisaged that an operator would manage the setup stage, then initiate the main observing program, and supervise the tear-down at the end. However, the main observing program would be managed by purpose-built software and be largely automatic.

Setup

The setup procedure would start with the operator in the OC allocating resources to an observing session. The operator might be starting an EOP session in which case they would choose an array of available antennas that have already been through system checks by local staff (pointing, SEFD, coherence checks etc) and made available to the pool. The next step would be to select the correlator or correlators that can be used to process the data. Local staff at the correlation centres would prepare their hardware and software, then make information on available resources (e.g. number of nodes) available to the OC so the correlator can be added to the correlator pool. The operator would choose other requirements, such as the desired scan SNR level, maximum number of sub arrays, data rates, default source catalogue etc.

The next step in the setup procedure is to make an observation of a bright source to obtain an initial fringe solution. The clock delays and rates obtained in this way are used in the main part of the observations. This fringe check also acts as a final array check. Absence of fringes to any antenna are indicative of a problem at the station that requires resolution before the antenna can join the array. The measured amplitude of the source on the various baselines can be used to estimate the SEFD of each antenna and verify that it meets with expectations. If fringes were weaker than expected or not found to a particular antenna, the operator would need to choose whether to leave the antenna out of the array until it is fixed, or to revise the antenna SEFD in the appropriate catalogue before proceeding.

Observation

Once the setup procedure has been completed, the telescopes and correlators are ready and the observations can begin. The operator would hand over operation of the array to the Observation Management Software (OMS). Figure 1 shows a flow chart that describes a proposed main observing cycle.
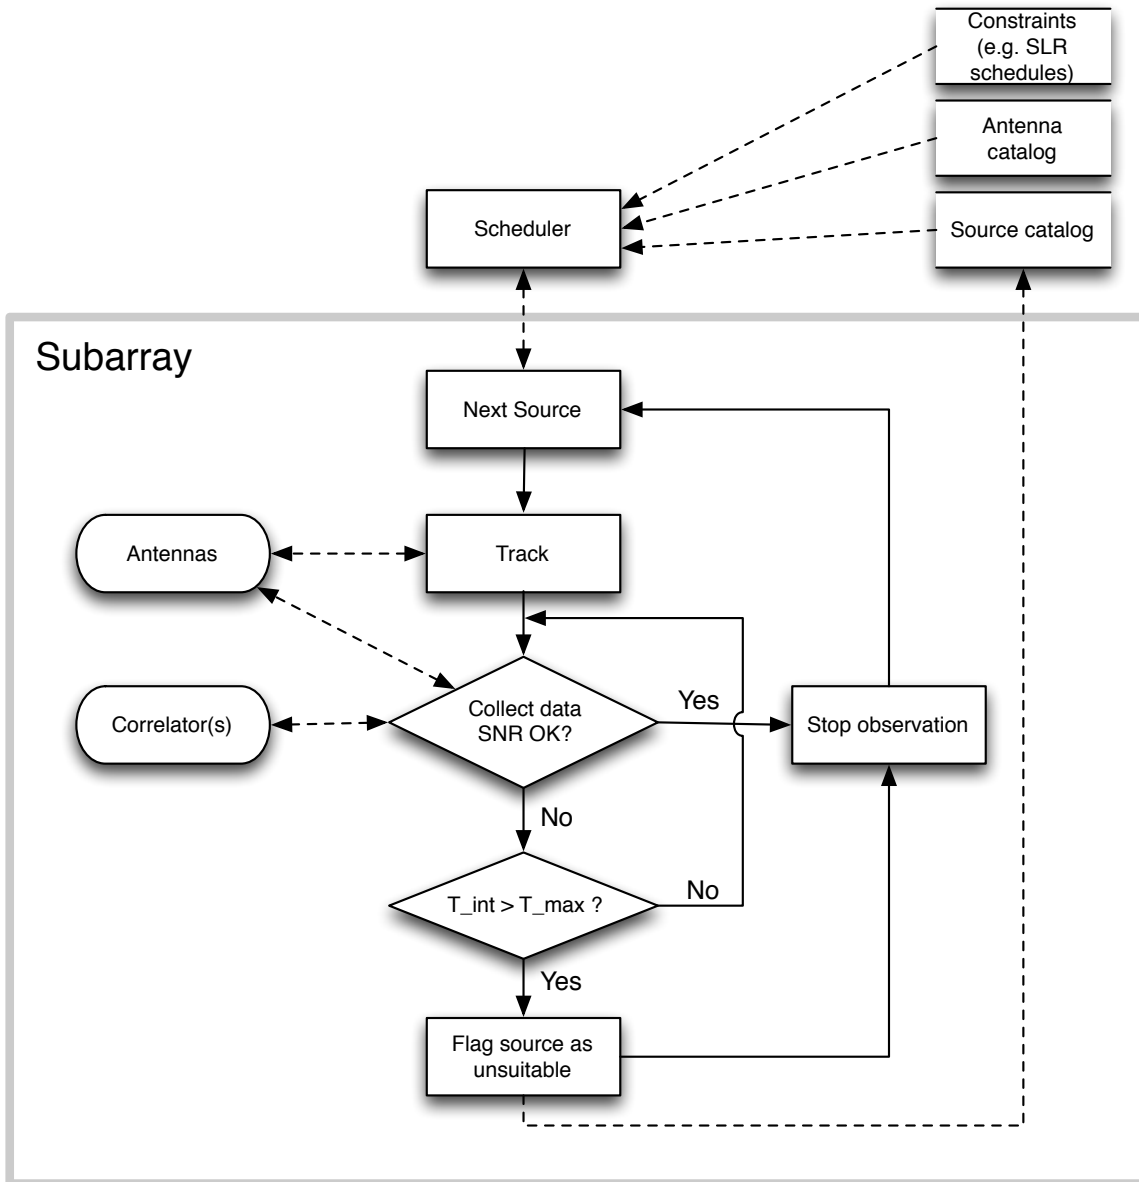

Figure 1: A flow chart describing the main observing cycle. Solid lines indicate flow while dashed lines indicate communication between elements.

At the heart of the OMS is the Scheduler. Much like the existing scheduling software (SKED or the VieVS scheduling module), the Scheduler would choose the next source to observe given a number of constraints such as recently observed sky positions, minimum distance to move, SNR and integration time limits, the length of the experiment etc. The Scheduler would draw on a catalogue of suitable sources, a list of available antennas and correlators, all updated during the observations. Other constraints could also be accepted such as the observing schedule for a co-located SLR at a site in order to avoid pointing an antenna toward an aircraft avoidance radar transmitter. The Scheduler would allocate a source and correlator to each subarray and the OMS would receive the information and instruct the antennas to move. Once the OMS has been notified by the antennas that they are on-source, an instruction would be sent to begin collecting data and the correlator

would be instructed to commence processing. Data collection and correlation would continue until the desired SNR is reached or until a pre-specified maximum integration time had been exceeded. The key here is to make the observation as short as possible but to reach a level of sensitivity that will result in a useful baseline solution. Therefore, real-time feedback from the correlator during the observation is essential.

If the maximum specified integration time is exceeded then this indicates either an unsuitable source or a problematic antenna. The OMS would make a choice to exclude an antenna, revise its SEFD or exclude a source as a result. Once the source observation is complete, the OMS would interrogate the Scheduler for a new source, correlator and subarray and continue observing.

2. Technology and software

A dynamic observation requires the following key components:
- Remotely controllable antennas
- Remotely controllable correlators
- Automated scheduling software
- OMS to coordinate the observations
- High capacity networks between antennas and correlators
The following paragraphs describe the current status of these elements and the work required to make them ready for dynamic observing.

Antenna control
In recent years, a significant amount of work has gone into remote control and monitoring of antennas with the eRemoteCtrl project (Neidhardt et al., 2010). This software is now used routinely for many IVS stations. For example, all three of the AuScope VLBI Array telescopes are controlled remotely using eRemoteCtrl as part of routine operations (Lovell et al., 2013). At the moment, all instructions to the antennas, samplers and recorders are sent either manually by an operator or via a pre-prepared schedule file. However, an interface to eRemoteCtrl or to it's associated server software on the host PC Field System computer could quite simply be built to allow instructions to be sent from the OMS.

Correlator control
Software correlators are now used in production geodetic VLBI and can be modified relatively easily to meet new requirements. Some of them already support eVLBI observations but additions or modifications to the software and the interface will be required to support the observing mode described here. In the case of the DiFX software correlator (Deller et al., 2011), a client/server communications layer is required to allow interaction with the OMS.

Also, DiFX currently expects an input VEX format schedule file and a v2d file with information such as station clock offsets and rates and predicted EOPs. This information would need to be provided and processed dynamically. DiFX would also require an eVLBI mode where a scan on a particular source could be stopped prematurely, triggered by feedback on SNR for example. This in turn requires software to monitor in real time the correlator output. Much of this already exists in DiFX through the difx_monitor software but additions would be required to trigger early integration stops and communicate back to the OMS.

Scheduling
Existing scheduling software such as SKED is already capable of automation. For example, given input parameters such as antennas, start and stop times, SNR targets, sky coverage requirements, source structure information etc, the software can prepare an optimised schedule file for an observing session.

In some ways, programs like SKED are already well suited to the task. In an automated scheduling run, the next source is chosen based on what has previously been scheduled given the pre-specified requirements and constraints. This is precisely what is required for dynamic observing except one source is chosen at a time based on the success (or failure) of previous observations, rather than filling an entire 24 h. The new challenges for scheduling software in a dynamic observing scenario are dealing with the changing input source list, antenna parameters, available antennas and allocating subarrays to correlators (which may require changes to the VEX format).

As with the correlator, a client/server layer is also required for the scheduling software to allow communication with the OMS.

The OMS
The Observation Management Software (OMS) would coordinate the observations by accepting some initial parameters and constraints (e.g. experiment name and type, start and stop times, antenna and correlator resources) and then conducting the observations through communications with the antennas, correlators and scheduling software. The design, building and testing of the OMS is likely to be a significant job given the need to manage several different types of experiment (simultaneously in some cases) with a high level of reliability and robustness.

High Capacity Networks
The high data recording rates required for VGOS will in turn require high capacity networks from the stations to the correlators to allow real-time data transfer and processing. This is discussed in more detail in the next section and is probably one of the more challenging requirements of VGOS as it spans national and international borders.

Development of Software and Operations Procedures
Probably the best approach in developing these new capabilities is to test the software and procedures at an IVS network station with two or more antennas, a software correlator and a fast local network. This would eliminate the additional complications of managing long-distance network connections and allow developers to concentrate on building and debugging software.

3. The transition to dynamic observations

There will inevitably be a period of transition from legacy systems to a fully capable VGOS network. During this time, some stations will have S/X systems and data recording rates of ~1 Gbps or less and others will have broadband VGOS systems and data rates of up to 32 Gbps. Some stations will have high capacity network connections capable of supporting the VGOS data rates in real time while others will have VGOS recording capability but insufficient network capacity. Further, the correlation facilities must have a factor of N times the network bandwidth capability (where N is the number of stations to correlate) in order to achieve real-time processing. Unless all stations can transfer all data to the correlator in real time, data will need to be kept either at the station or the correlator until they can all be

brought together for processing. Even though these network bottlenecks prevent a full VGOS implementation in the medium term, the possibility still exists to implement the most important aspects of dynamic observing with only modest network capacities.

Existing software on Mark5 recorders such as jive5ab already permit simultaneous data streaming and local recording, so it is possible to send data for real-time processing for a subset of stations[1]. Further, if a station or correlator has a limited connection, it may be possible to send a subset of a full data-stream (e.g. a single polarisation, 1-bit data instead of 2-bit, and/or a single or sub-band) to the correlator for real-time analysis which would be sufficient for SNR measurement, dynamic scheduling and initial data products within 24 h, with the complete data set transferred later. If multiple correlators are available but the network bottlenecks are into the correlators rather than out of the stations, the data streams could be split by IF and sent to different correlators simultaneously. It should be noted that software development and/or testing is still required for simultaneous recording and data transfer on new recoding systems such as Mark6 and FlexBuff.

Pre-experiment checks such as fringe-finding can of course be carried out in non-real-time mode without a serious impact on observations. For example, a 10 s integration on a bright source at 32 Gbps recording rate would take a little over 5 minutes to transfer to a correlator with a 1 Gbps connection to the station. Even centrally coordinated pre-experiment checks, which would serve to check for telescope readiness and performance and allow the scheduling software to update antenna SEFDs and thus on-source time calculations, would be a significant step forward.

As total AGN flux density variability on timescales of months to years occurs in the sub-milliarcsec-scale jet, regular input from single-dish or short-baseline interferometer monitoring programs could also be used to keep flux density databases up to date. A flux density monitoring program could be coordinated between IVS stations to achieve this. Imaging of sources every few weeks from IVS data may be sufficient to determine the suitability of sources and therefore reduce the need for real-time SNR determination to optimise scan times. However, a more detailed study to assess the relative merits of these approaches is probably required.

Acknowledgements

References

Deller, A. T., Brisken,W. F., Phillips, C. J., Morgan, J., Alef, W., Cappallo, R., Middelberg, E., Romney, J., Rottmann, H., Tingay, S. J., and Wayth, R. (2011). DiFX-2: A More Flexible, Efficient, Robust, and Powerful Software Correlator. PASP , 123 , 275–287.

Lovell, J.E.J. et al (2013). The AuScope geodetic VLBI array. Journal of Geodesy, 87, 527-538

Neidhardt, A., Ettl, M., Pl¨otz, C., M¨uhlbauer, M., Hase, H., Sobarzo Guzm´an, S., Herrera Ruztort, C., Alef, W., Rottmann, H., and Himwich, E. (2010). Interacting with radio telescopes in real-time during VLBI sessions using e-control. In Proceedings of the 10th European VLBI Network Symposium and EVN Users Meeting: VLBI and the new generation of radio arrays. September 20-24, 2010. Manchester, UK

---

[1] Note that the older Streamstor cards, mostly found in Mark5A recorders, can manage transfer rates of up to 512 Mbps. The newer AMAZON-based Mark5 units can support rates of up to 1 Gbps.