

VGOS Data Transmission and Correlation Plan

Version: November 18, 2014

Bill Petrachenko, Alessandra Bertarini, Walter Alef, Dirk Behrend, Roger Cappallo, Hayo Hase, Chopo Ma, Arthur Niell, Axel Nothnagel, Xiuzhong Zhang

Abstract. The establishment of the VLBI Global Observing System (VGOS) requires a projection of the needs for data transmission as well as for correlation to achieve the best possible results and a smooth implementation phase from the very beginning. In this document we outline the requirements and possible developments, which are related to (a) the transport of the raw VLBI data from the stations to the correlators and (b) the correlation and fringe fitting process itself.

1. Introduction

The implementation of the VLBI Global Observing System (VGOS) is being guided by the VGOS Program Executive Group (VPEG), which documents its discussions and considerations in a series of VGOS plans. Following the *VGOS Observation Plan* (Petrachenko et al. 2014), this document now addresses the requirements for VGOS data transmission and correlation. In technical terms, a fundamental assumption for this plan is that the observation cycle of one source has a duration of 30s including a typical on-source integration time of 7.5 s. Data is acquired at 16 Gbps. Assuming that the sequence of observations can be realized as set forth in the *VGOS Observation Plan*, the VGOS data production rates increase from 58 TeraByte (TB) per day in the year 2015 to 1037 TB/day in 2020 (Tab. 1). The last column, “data/day at correlator”, assumes that there is only one monolithic correlator in use. With more than one correlator, sessions would be distributed such that, on average, the total correlator load would be divided in proportion to the available capacity at each correlator. Although interesting from the point of view of data transmission and cost sharing, a highly distributed correlator is not discussed here.

Year	# of sites	hours of obs/day	data/day/site (TB)	data/day at correlator (TB)
2015	8	4	7.2	58
2016	10	8	14.4	144
2017	16	8	14.4	230
2018	20	10	18.0	360
2019	24	12	21.6	518
2020	24	24	43.2	1037

Table 1. Estimates of VGOS data production rates for the years 2015 to 2020.

This document will attempt to predict VGOS requirements for e-transfer, disk pack shipping, media requirements per station, number of playback units at the correlator, correlator cores, internal correlator network, and on-line correlator memory for the years 2015 to 2020. In this version of the document, the increased data transmission and correlation load related to multi-antenna sites has not been taken into account. For twin (VGOS) telescope sites this could lead to as much as a doubling of the data production rate if the most aggressive observing scenarios are assumed for both antennas. For sites that will continue to operate a legacy antenna along with the new VGOS antenna(s), the added data production rate will typically be significantly less severe than for twin telescope sites due to the slower slew speed of the legacy antenna. Furthermore no attempt will be made here to estimate power requirements for the correlator although this could become an important aspect of correlator operational costs; hence it may be considered in a later version of this document. It is assumed that the correlator will be a software correlator and that its architecture will be similar to existing software correlators.

2. e-Transfer Rates

Based on the information in Table 1 it is possible to calculate requirements for sustained data transmission rate from each site and for a cumulative sustained reception rate at the correlator. These are calculated simply as the total daily data production divided by the number of seconds in a day. The cumulative sustained reception rate at the correlator in column five of the table assumes that all stations transmit data to the correlator via e-transfer and hence is the maximum value expected. When specifying a requirement for network data rate it is common to account for network overhead by applying about a 40% margin above the sustained data rate to arrive at the required network data rate. Stations planning to transfer data over the Internet have to increase their capacity from 1.0 Gigabit per second (Gbps) in 2015 to 5.6 Gbps in 2020 (Tab. 2). If the data of all stations would be transferred over the Internet, the correlator(s) need an increase in network capacity from 8 Gbps to 134 Gbps. For some stations and correlators where direct inexpensive access to a high speed network is not practical, it may be cost effective and efficient to physically transport modules to or from a nearby node on a network trunk and then e-transfer from that point onward.

Year	data rate at each site (Gbps)	network data rate at each site (Gbps)	data rate at correlator (Gbps)	network data rate at correlator (Gbps)
2015	0.7	1.0	5	8
2016	1.3	1.9	13	19
2017	1.3	1.9	21	30
2018	1.7	2.4	34	48
2019	2.0	2.8	48	68
2020	4.0	5.6	96	134

Table 2. Estimates of sustained e-transfer rates and network data rates.

3. Record Media Shipment

The other alternative for data transmission is to physically ship record media. By 2015 it is expected that new 8-pack recording modules will be based on 6 TB disks (i.e. 48 TB total for each 8-pack module). In Table 3, the data production rates from Table 1 are re-expressed in terms of the number of 8-pack record modules required per day. In most cases disk packs will be significantly less than full at the end of 24 hours.

Year	8-pack module capacity (TB)	data produced per day (# of 8-pack modules)
2015	48	0.15
2016	48	0.30
2017	48	0.30
2018	48	0.37
2019	48	0.45
2020	48	0.90

Table 3. Estimates of the number of 8-pack record modules required per day at each station.

4. Media Requirements per Station

There are three cases that need to be considered for estimating a stations media requirements:

- *All stations ship disk packs.* Any station that ships disk packs needs enough media to handle about eight weeks of recording (one week for shipping to the station, one week for shipping to the correlator, two weeks for correlator schedule slippages and post-processing, and four weeks to reduce costs at the correlator through bulk monthly shipment of media to the stations);
- *Some stations ship disk packs and some stations e-transfer data.* In this case, the stations that ship disk packs still need eight weeks of media; but the stations that e-transfer data only need enough media to handle three weeks of recording before modules can be erased for re-use (one week while the shipped disks travel to the correlator and two weeks for correlator schedule slippages and post-processing) [An alternative for stations that e-transfer is to have enough on-line RAID storage to store the raw data for three weeks.];
- *All stations e-transfer data.* For a network where all stations e-transfer data there are two cases:
 - *Non real time.* In this case each station needs two weeks of media to account for correlator schedule and data transmission slippages, and to allow for completion of post-processing

- *Real time.* In this case correlator schedule slippages are not tolerated and only limited buffering is required to account for network overhead and unexpected network outages.

In Table 4, the media requirements per station are shown. As in the previous section, it is assumed that 48 TB modules will be used. Even though some stations will have more than adequate network capability, it is recommended that they still maintain a capability for removable data recording media since correlators may not be able to handle the required network traffic especially during early operations. It is further recommended that the IVS decide on a compatibility standard for record media so that correlators will not need to support more than one type of media.

Year	module size (Tbytes)	Data per day (# of modules)	Days/module (modules/year 1-way shipped)	Sites that ship data (modules needed/site)	Sites that e-transfer in a mixed network (RAID TB)	All sites e-transfer - not real time (RAID TB)
2015	48	0.15	6 (21)	7	4 (151)	3 (101)
2016	48	0.30	3 (122)	19	7 (302)	5 (202)
2017	48	0.30	3 (122)	19	7 (302)	5 (202)
2018	48	0.37	2 (183)	28	11 (373)	7 (249)
2019	48	0.45	2 (183)	28	11 (453)	7 (303)
2020	48	0.90	1 (365)	56	21 (907)	14 (605)

Table 4. Media requirements per station assuming the case where data are shipped only when media is nearly full. As an alternative to removable disk modules for e-transfer data, the numbers in brackets in the last two columns of the table represent the amount of on-line RAID storage (TB) that would be needed to store all e-transferred data from a single station until it can be released.

5. Playback systems at the correlator

It is assumed that the number of playback units at the correlator will be large enough to make possible a semi-automated operating mode where recorded modules never need to be exchanged more often than once per day. In this section, Mk6 units are used as examples although other types of compatible and competitive playback units may also be developed in the course of time. It is assumed that a data acquisition duty cycle of less than 25% will be used so that no more than a single disk pack at a time needs to be mounted at a station and that each Mk6 playback unit at the correlator will be able to mount up to four recorded modules. There are three cases that need to be considered:

- *All stations ship data.* In this case three criteria need to be met:

- There needs to be a place for a recorded module from each station to be loaded. For example, since each Mk6 can handle four recorded modules, an eight station network would require at least two Mk6s.
- The cumulative sustained data rate of the correlator needs to be met. For example, since each Mk6 can play back at 16 Gbps (or somewhat more) when four disk packs are loaded, a sustained data rate of 21 Gbps would require at least two Mk6s.
- It must be possible to store the total volume of data acquired in 24 hours. For example, since each Mk6 fully loaded with four 48-TB record modules can store 192-TB of data, a 230 TB session would require at least two Mk6s.
- *All stations e-transfer data.* In this case only the latter two criteria need to be met but now an equal number of Mk6s is required to record data from the network and to play it back to the correlator so that the number of units is doubled. For example, for a sustained data rate of 21 Gbps and a total volume of 230 Tbytes, two Mk6s would be required to record from the network and two to play back to the correlator, for a total of 4 Mk6s.
- *Some stations ship data and some stations e-transfer data.* In this case, the number of Mk6s required will lie somewhere between the cases where all stations ship and all stations e-transfer. The final number will depend on the ratio of stations shipping and e-transferring.

See Table 5 for an estimate of the number of Mk6s required at the correlator.

Year	# of sites	Data rate at correlator (Gbps)	Data volume (TB)	# of Mk6s (all sites ship data)	# of Mk6s (all sites e-transfer data)
2015	8	5	58	2	2
2016	10	13	144	3	4
2017	16	21	230	4	4
2018	20	33	360	5	6
2019	24	48	518	6	8
2020	24	96	1037	6	16

Table 5. Estimate of number of Mk6s required at correlator.

6. Correlator Cores

In principle, it should be easy to estimate the VGOS correlator CPU requirements since the processing load of a DiFX correlator (or any software correlator with a similar architecture) increases more or less linearly with the number of stations in a scan. [With a typical DiFX correlator this linear relation is valid up to about 15 stations which is in line with the vast majority of scans expected for the VGOS network even up to 2020 and

beyond.] As a result, the number of CPU operations required to keep up with data acquisition should simply be proportional to the cumulative sustained input data rate of the correlator. (See Table 2.) Unfortunately, performance benchmarks that relate correlator input data rate to required number of correlator cores are often misleading when used to extrapolate to other correlator architectures, generations of technology, and observing scenarios. This is largely due to the fact that complex real time interactions amongst various hardware components and software processes lead to bottlenecks that move from location to location in the system and change in severity depending on correlator size, architecture, technology, and observing configuration. In the absence of something better to do, a universal scaling factor of about 40 correlator cores per Gbps was derived based on the processing of a single broadband session at the Haystack DiFX correlator and that scaling factor was then used to extrapolate to the VGOS correlator configurations shown in Table 6. For the reasons described above it is not expected that the estimates in the table will be accurate to better than about a factor of three in either direction.

Year	# of correlator cores
2015	200
2016	600
2017	900
2018	1400
2019	2000
2020	3900

Table 6. *Estimates of correlator cores required. It is expected that these values are not accurate to better than a factor of about three.*

7. Internal data rate at the correlator

As a lower limit, the internal data rate of the correlator needs to handle the sustained data acquisition rate. However, since data may at times need to be transmitted more than once before they arrive at their final destination for correlation, an efficiency factor needs to be applied to arrive at a realistic recommendation for internal data rate. It is difficult to be exact about the size of this factor since it will change both with correlator architecture and session configuration. However, for the purpose of this document, a factor of 40% is considered reasonable. If data is e-transferred, then double that rate is required so that transfers from the network to storage and storage to the correlator cores can be handled simultaneously (Tab. 7). Although current correlators are built around 10 Gbps network technology, modern equipment is increasingly able to handle higher data rates. It is recommended that a move to higher rate (e.g., 40 Gbps) technology be considered.

Year	Data rate at correlator (Gbps)	Internal data rate (all sites ship data) (Gbps)	Internal data rate (all sites e-transfer data) (Gbps)
2015	5	7	14
2016	13	19	37
2017	21	30	59
2018	34	48	96
2019	48	68	135
2020	96	135	269

Table 7. Estimates of required internal data rates at correlator.

8. Local storage at the correlator

In addition to the correlation playback units, each correlator needs sufficient local storage capacity for buffering data. Taking into account the data streams coming in by e-transfer and a correlation backlog of two observing days, correlators need local storage capacity of at least 500 TB. For this, either commercial RAID systems or dedicated VLBI units are available. For the dedicated VLBI units, the current MK6 units are used as examples. Along with high storage density and data rates they also have the advantage of removable media.

9. Current and Planned IVS Correlator Resources

Location	Correlator Cores		Mk6s		External Network (Gbps)	
	Now	Planned	Now	Planned	Now	Planned
Bonn	488	1000-1500	2	6 (2 slot)	1	no plan
USNO	512	1024	0	?	1	10
Haystack	100	~300	3	6	20	no plan
Shanghai	64	1000	2	?	1	no plan
Tsukuba	92	256	(24.9 TB)	(513 TB)	10	no plan

Table 8. Current (2014) and planned correlator resources.

Table 8 shows the current and planned resources available for IVS processing at existing software correlators. In most cases these correlators will be shared with other astronomical and/or space programs so only a fraction of the capability will be available for geodesy. Although other aspects of the correlator are important, e.g. internal network resources and internal storage, the table only considers correlator cores, Mk6s and the external network connection. In the case of the Tsukuba correlator, there are no plans for

Mk6s so internal storage is displayed and it should be noted that their planned upgrade has in fact recently been implemented. Correlators that might be of interest but are not included in the table are the VLBA correlator which is used for processing RDV sessions, the Curtin correlator used for AUSTRAL sessions, the Kashima correlator and the Russian correlator. Their use for VGOS sessions will have to be discussed. Available correlator resources in Table 8 can be compared with projected correlator requirements in Tables 6, 5 and 2 respectively.

Station	Network Data Rate(Gbps)
Noto	10
Kokee	0.1
Westford	20
GGAO	1
Ishioka	10
Sheshan	1
Wettzell	1
Yebees	10

Table 9. Network data rates at sites most likely to be ready for broadband observing in 2015.

Table 9 shows network data rates currently available at the VGOS stations most likely to be ready for the broadband test campaigns in 2015. In general data rates vary considerably from site to site. These values should be compared with the projected requirements in Table 2.

10. Recommendations

- Correlator Cores.** Within the next year or two there could be significantly more than 3000 correlator cores available at IVS correlators. According to Table 6 this would satisfy VGOS correlator core requirements up to about 2019. However, this takes into consideration neither the fact that most IVS correlators are shared with other applications (including legacy S/X observing) nor the fact that the estimates in Table 6 are only reliable to within a factor of about three. Furthermore, since most correlators are currently being upgraded with respect to correlator cores it is unlikely that sponsors will fund further upgrades in the near future. As a result it is recommended that the current set of upgrades be completed and that the situation with respect to correlator cores be reviewed in about two years. Over that time, broadband data will have routinely been processed using modern correlator cores so that better estimates of the relation between data throughput and correlator cores can be made. At the same time, more accurate schedules for the roll out of VGOS stations and 24/7 operation will be known making predictions of correlator core requirements more realistic.

- **Dedicated VLBI disk units.** At the moment, no VLBI disk unit with competing capacity to the MK6 units is available. For this reason, MK6 units are used as examples. At present only a few Mk6s units are deployed at observing sites and at correlators. It is recommended that correlators aim to have at least four Mk6s (or equivalent) on the short term and as many as eight on the longer term.
- **External network connections at correlator.** To achieve high precision low latency results, network connections at correlators should be upgraded to at least 10 Gbps on the short term with significantly higher rates (40 Gbps) required in the future to support 24/7 operations with the full network.
- **Network connections at stations.** If real time processing is a priority, network VGOS stations should be upgraded to handle at least 2 Gbps on the short term with 5.6 Gbps required for full 24/7 operations.
- Even though some stations will have more than adequate network capability, it is recommended that they still maintain a capability for removable data recording media since correlators may not be able to handle the required network traffic during early operations.
- It is recommended that the IVS decide on a compatibility standard for record media so that correlators will not need to support more than one type of media.
- **Internal network technology at the correlator.** Although current correlators are built around 10 Gbps network technology, modern equipment is increasingly able to handle higher data rates. It is recommended that a move to higher rate (e.g., 40 Gbps) technology be considered.

References

Petrachenko B, Behrend D, Hase H, Ma C, Niell A, Nothnagel A, Zhang X (2014): **VGOS Observing Plan** (February 13, 2014) http://ivscc.gsfc.nasa.gov/technology/vgos-docs/vgos_observing_plan_140213.pdf