Australian VLBI Correlation Center

Tiege McCarthy[1], Jamie McCallum[1], Lucia McCallum[1], Oleg Titov[2]

Abstract We outline the progress and current plan for development of an Australian VLBI Correlation Center operated by the University of Tasmania. This center will fill the gap in Australian VLBI capabilities and initially focus on the correlation of geodetic VLBI data from Australia and the Asia-Pacific region, with the ultimate goal of becoming recognized by the IVS as an international VLBI Correlation Center. The center will be designed from the ground up with the data volumes of VGOS in mind and aims to address the current lack of Correlation Centers that are equipped to handle the routine correlation of VGOS sessions. We aim to transition from a local correlation cluster to an HPC environment, allowing us to overcome data I/O issues that are inherent to the high data rates of VGOS. This center will allow for better utilization of the Australian VLBI network, increasing cadence of observations with shorter turnaround times on processing.

Keywords Correlation, VLBI, VGOS

1 Introduction

In order to realize the goals of the VLBI Global Observing System (VGOS), data logistics and processing need to be considered and scaled up accordingly. High experiment cadence combined with higher data rates, resulting in larger data volumes, will push existing infrastructure beyond its limits. The ability to access and read data rapidly is a critical component of the data correlation process, and the greater than an order of magnitude increase in data rate for VGOS experiments is going to provide a significant challenge in the future.

The IVS Infrastructure Development Plan 2030 [1] identifies the need for more Correlation Centers as we move into the VGOS era. In partnership with Geoscience Australia, the University of Tasmania (UTAS) is establishing an Australian VLBI Correlation Center which aims to provide the missing piece of Australian geodetic VLBI capabilities. Development of expertise, documentation, and facilities will move us from our current ad-hoc correlation support to a consistent routine correlation service for the Asia-Pacific region and ultimately the IVS as a whole.

In this paper we outline our existing correlation facilities and workload, along with our plans for this future Australian VLBI Correlation Center.

2 Current UTAS VLBI Infrastructure

UTAS operates a geodetic VLBI array situated across Australia, with two antennas located at the Mount Pleasant Observatory (Hb and Ho) in southern Tasmania, one station in Katherine (Ke, Northern Territory) and one station in Yarragadee (Yg, Western Australia). The remote stations in Katherine and Yarragadee currently have poor network connections and therefore require experiment data to be shipped on disksets to Mount Pleasant Observatory for processing. Planned upgrades to the network connections for the remote stations will hopefully remove the need for the shipping of disksets in the not-too-distant future.

1. School of Natural Sciences, The University of Tasmania, Private Bag 37, Hobart, Australia
2. Geoscience Australia, PO Box 378, Canberra, Australia

2.1 Current Correlation Responsibilities

Currently UTAS is responsible for the correlation of the following sessions:

- Approximately two mixed-mode (legacy S/X stations combined with VGOS stations in a pseudo-S/X observing mode) sessions, rotating fortnightly between the AUA and AUM code sessions. These sessions are typically 24-hour experiments with 4–6 stations total.
- A weekly southern Intensive mixed-mode session with two-station short turnaround processing followed by a later three-station processing.
- VGOS test sessions between the VGOS-equipped Hobart 12-m and Katherine 12-m stations
- Occasional Asia-Oceania VLBI sessions. Approximately three times per year.

Our current correlation facilities are able to handle this existing workload; however, we currently do not have much room to expand our responsibilities. The establishment of this Australian Correlation Center will allow us to push our capabilties and provide a correlation service to a wider audience.

2.2 Mount Pleasant Correlation Cluster

Correlation at Mount Pleasant Observatory utilizes a small correlation cluster, along with Flexbuff data recording/storage machines. Disksets from remote stations are loaded into these machines so that they can be accessed alongside the locally recorded data. Each of the Flexbuff machines is connected to the observatory network switch with a 10-Gbps line.

The cluster itself is comprised of ten machines, that each have 20 available cores and 64 GB of memory. The machines from the cluster, similarly to the Flexbuff machines, are each connected with a 10-Gbps link to the network switch. Correlation is set up and performed on this cluster, with the subsequent fringe data sent to another local machine for fringe-fitting and post-processing before the vgosDB database is generated and distributed to the Analysis Centers. The correlation data and fringe-fitting data is then subsequently archived on a Research Data Storage Initiative (RDSI) node.

A typical five-station, 24-hour AUA session, with 9-Gbps total data rate when recording (the pseudo-S/X mode for VGOS stations is 3-Gbps data rate, legacy S/X is 1 Gbps), results in approximately 60 TB of data. We can currently correlate these sessions at approximately real time (24–30 hours). Comparing these numbers to a typical five-station VGOS session however, we are looking at five times higher data rates and data volumes, which will push this facility out of being able to comfortably perform routine correlation services.

3 Discussion of Solutions

The geographic location of UTAS (on the island state of Tasmania, off the southern coast of the Australian mainland), combined with existing telescope infrastructure and our increasing requirements provide a unique challenge when determining what the best course of action is for establishing this Correlation Center. The current solutions we are investigating include a substantial upgrade to our current in-house correlation setup or moving toward high performance computing (HPC) solutions.

3.1 Upgrade Existing Correlation Facilities

Upgrades to our existing correlation infrastructure can be made in three main ways:

1. Processing power;
2. Data input/output;
3. Data storage.

Processing power is a relatively cheap and easy to upgrade; however, it is unlikely to provide a significant upgrade to correlator performance until our correlation workloads move toward much higher station count experiments. When looking at core count versus correlation time for a fixed number of stations, in all cases we see great scaling with increasing core count up to a break point where we hit extreme diminishing returns (see Figure 1 for a three-station example using some southern Intensive data).

Data input/output on the other hand is where we believe the biggest improvements can be made; however, this is still an area we are actively investigating.

Additionally, when compared to upgrading processing power, data I/O improvements are much more expensive and difficult to implement into our existing correlation cluster.

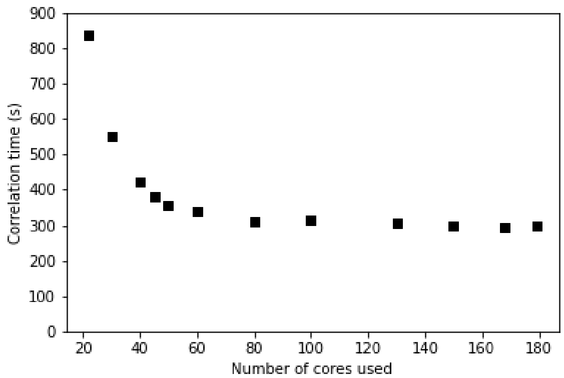

Fig. 1 Correlation time versus core count scaling for our current Mount Pleasant Observatory correlation cluster, using ten scans (580 seconds) of a three-station mixed-mode southern Intensive session.

Another problem with upgrading our current facility is that we are still limited by our network link across Bass Strait, the body of water between Tasmania and mainland Australia. Currently we share a 10-Gbps link across to the mainland with many other research facilities in Tasmania—meaning e-transfer data rates to our current facility is somewhat variable. This bottleneck may become even more pronounced, as our remote stations get upgraded with fiber connections and e-transferring data from these sites becomes standard practice.

3.2 Move to HPC Solution

We have access to two different HPC facilities as part of this project: (1) the UTAS run Tasmanian Partnership for Advanced Computing (TPAC) *kunanyi* HPC in Tasmania and (2) the National Computational Infrastructure (NCI) *Gadi* HPC in the Australian Capital Territory.

3.2.1 *kunanyi*

The TPAC *kunanyi* HPC has 6,720 cores distributed across 240 nodes with 128 GB of memory available on each node. The HPC has 1 PB of scratch storage, PBS job scheduling, and high-data rate (InfiniBand) connectivity between the nodes. This hardware should completely satisfy all performance and I/O needs for the foreseeable future.

However, moving to an HPC does not solve the issues of data transport and logistics, and getting large volumes of data onto *kunanyi* is a currently unsolved challenge. Data physically shipped to UTAS will still need to be moved onto the HPC. Theoretically, the Mount Pleasant Observatory is connected to *kunanyi* via a 10-Gbps link; however, current transfer testing has been hitting a 1-Gbps limit, which we hope will be rectified in the near future. One benefit of this HPC being local is the potential for a local data ingest station, where we can load Flexbuff disk sets directly into accessible storage on site. The feasibility of this is currently being investigated. On the other hand, the HPC being local means we suffer from the same reliance on our shared 10-Gbps link to mainland Australia.

3.2.2 *Gadi*

NCI's *Gadi* has 3,074 nodes each with 48 cores (147,552 total), 192 GB of memory, and similar to *kunanyi* high-data rate (InfiniBand) inter-node connectivity. It uses PBS job scheduling and will more than satisfy our processing and data I/O needs. Access to *Gadi* is provided through our partnership with Geoscience Australia, which are a major collaborator with NCI.

Again, similar to *kunanyi*, data logistics will also be an issue for *Gadi*. Currently, despite our 10-Gbps line to the mainland, our typical realistic transfer rate onto *Gadi* is between 2–3 Gbps. This makes shipping multiple stations worth of data to Hobart then subsequently transferring onto *Gadi* not feasible, even transferring VGOS data for just the Hobart 12-m will take 48 hours for a 24-hour VGOS session. However, we do have the possibility for a data-relay station at the Canberra Geoscience Australia (GA) headquarters, which we can ship data to from the remote stations and then transfer from GA onto Gadi at approximately 10 Gbps. This would alleviate the issue of 'double transferring' and significantly cut down transfer times. This relay station would only be necessary as a stopgap solution until fiber connections to our remote stations are online in the next couple of years.

Gadi as a mainland based solution is very attractive moving forward and will scale better with future infrastructure upgrades. It also removes the Bass Strait network connection as a potential bottleneck, as we will only need to transfer data from one VGOS station across it (Hb) rather than potentially multiple remote stations.

3.3 Distributing Correlation across Infrastructure

Because moving the data around is our limiting factor for all solutions, potentially distributing our correlation jobs around all the potential solutions makes the most sense, especially for the short term while we are limited by physically shipping data. Depending on the experiment, data can be shipped to various locations—multiple large correlation jobs can be distributed to different correlation hardware, while our local cluster continues to handle smaller jobs or legacy/mixed-mode correlation.

While this solution doesn't drastically decrease processing latency, it could provide quite high overall throughput of the center. Additionally, as connectivity increases to our remote stations, we can be even more flexible with how we distribute correlation workloads.

A distributed correlation solution such as this carries with it the primary downside of drastically increased bookkeeping required to keep on top of the data logistics.

4 Summary

We aim to establish an Australian VLBI Correlation Center that will provide VGOS capable correlation services to the Asia-Pacific Region and the IVS as a whole. We outline our current existing correlation infrastructure and correlation workload, along with the three primary solutions we are investigating to provide suitable hardware moving forward. From our initial investigation, the two HPC solutions appear to be the most promising, with our existing cluster remaining for small projects.

We also discuss ways to navigate the current limiting factor of data transfer from remote stations by distributing our correlation workload across multiple facilities in order to increase total correlation throughput.

Acknowledgements

This work is funded by Geoscience Australia as part of the Positioning Australia project.

References

1. A. Nothnagel, J. Anderson, D. Behrend, et al., "IVS Infrastructure Development Plan 2030." In: International VLBI Service for Geodesy and Astrometry 2019+2020 Biennial Report, edited by D. Behrend, K. L. Armstrong, and K. D. Baver, NASA/TP-20210021389, pp. 3–11, 2021.