MIT Haystack Observatory VLBI Correlator: Biennial Report

Mike Titus, John Barrett, Roger Cappallo, Brian Corey, Pedro Elosegui, Arthur Niell, Chester Ruszczyk, Jason SooHoo

Abstract This report summarizes the activities of the VLBI correlator at the Massachusetts Institute of Technology (MIT) Haystack Observatory during the period 2017–2018.

1 Introduction

The distributed FX-type (DiFX) VLBI software correlator (Deller et al., 2011) of the MIT Haystack Observatory, located in Westford, Massachusetts, is supported by the NASA Space Geodesy Program (SGP) and the National Science Foundation (NSF). It is dedicated approximately equally to the geodetic pursuits for the International VLBI Service (IVS) and to radio astronomy imaging for the Event Horizon Telescope (EHT) project. The MIT Haystack correlator serves as a development system for testing new correlation modes, such as those needed for observations with the next-generation VLBI system — the so-called VLBI Global Observing System (VGOS) — and for recorder developments, such as the Mark6 system. Software support is also provided to similar DiFX installations at the U.S. Naval Observatory (USNO) and the Max Planck Institut für Radioastronomy in Bonn (MPIfR), Germany, and to the general IVS community for processing of VLBI experiments.

MIT Haystack Observatory

MIT Haystack Correlator

IVS 2017+2018 Biennial Report

2 Summary of Activities

2.1 General VGOS Activities

The last two years have seen a significant expansion and ramping up of regular VGOS observations. Progress has mostly moved beyond a period of basic debugging of fundamental problems, as described in the last report, to observing a regular series of VGOS experiments once every two weeks with six stations.

During this period the VGOS stations at Wettzell, Yebes, and ("northeast") Onsala have joined the US stations (i.e., Westford, GGAO, and KPGO) on a regular basis. Occasionally, ("southwest") Onsala and Ishioka have also joined the group, making for a VGOS network of eight stations. There is still a moderate level of problem diagnosis and repair embedded in the process, but in general all the participating stations fairly routinely generate data worthy of inclusion in a database for further data processing and analysis.

In addition, a VGOS network of six stations successfully participated in the Continuous VLBI Campaign 2017 (CONT17) for five continuous days, 4–8 December 2017. The broadband VGOS network ran simultaneously with two legacy VLBI networks, which only observed with the traditional dual-frequency (S/X) bands for a period of 15 days, from 28 November to 12 December 2017.

As a measure of the maturing of the VGOS network, databases from the regular, bi-weekly VGOS observations are now being publicly distributed.

2.2 CONT17 Campaign

Significant effort went into preparatory observations, execution, and post-processing of the CONT17 campaign. A pilot pre-CONT17 campaign spanning two consecutive days was conducted in preparation. There was much preparatory work in station problem diagnosis and repair and in post-processing problem discovery, diagnosis, and mitigation. For example, correct VGOS-specific clock and sampler delay setting techniques became a major issue in VGOS processing, but, ultimately, results from VGOS and the legacy network were well aligned.

2.3 Tool Development for VGOS Processing Automation

We have developed software tools and an accompanying users manual to describe and facilitate the correlation and post-processing steps of VGOS observing (Barrett et al., 2019), as this process is significantly different from traditional S/X observing. This effort is a significant knowledge-transfer exercise of VGOS processing capabilities to other correlators and to the IVS.

2.4 VGOS Equipment Verification and Validation

The MIT broadband signal chain for the new VGOS station at McDonald Observatory (MGO; Merkowitz et al., 2018) has undergone comprehensive testing. This includes parallel-rack testing whereby two VGOS backend systems, the Westford and MGO backends, ran in zero-baseline mode at the Westford station during regularly scheduled VGOS sessions. Both systems were correlated at MIT Haystack for verification and validation of the MGO signal chain.

2.5 Digital Backend Testing

Tests of non-standard modes of observing with RDBEs, such as full 512 MHz bandwidth with 16 Gb/sec recording using all RDBE channels and testing a proposed new frequency sequence (B. Petrachenko, private communications), have been conducted and are being analyzed. Tests of VGOS-compliant (1024 MHz bandwidth) R2DBEs have been performed during this period, validating their use for VGOS observations at the Westford site.

2.6 Mixed-mode Observations

Results from a mixed-mode (using both legacy and VGOS antennas) test session have been obtained, and two more mixed-mode sessions were observed since the last report. The RD1606 mixed-mode experiment with Westford tagging along was shown to fully validate the technique and was released to the IVS. Since that time, two more mixed-mode sessions have been observed, first with Westford and GGAO recording simultaneously in mixed mode in RD1804, and then with KPGO joining in for RD1810. These experiments are in post-production and production, respectively.

Parallel development of the MIT Haystack Observatory Post-processing System (HOPS) software package has continued apace in order to accommodate this highly non-standard mode which is incompatible with current norms. Most changes have been directed at simplifying the data processing pipeline. This is necessary as the current pipeline is not sustainable for bulk production.

2.7 Cluster Upgrade

A major expansion of the correlator cluster was implemented in 2017 under the auspices of the EHT project. Thirty-eight new compute servers were purchased and installed. Together they host 2x10 core (Intel Xeon) CPUs, three 100 Gb/sec Ethernet backbone switches, and (Mellinox) network cards supporting a Gb/sec Ethernet fabric with supporting racks and power distribution. The upgrade has greatly increased the bandwidth and computational power of the correlation backbone infrastructure. Additionally, four new Mark6 playback units have been added, increasing the total number of available playback units to eleven. Approximately 38 TB of storage space has also been added for data.

In parallel, global IT management tools were implemented in order to make configuration and control of this new large cluster of machines manageable. Since the last report, all of the equipment from the geodetic and EHT clusters has been combined into one uniform cluster. The geodetic (Infiniband) fabric was retired in favor of the Gb Ethernet fabric. As a result, the geodetic projects that include VGOS antennas have benefited from the upgrade in having access to this hardware for their correlation tasks.

2.8 Clock-summit Meetings

A teleconference series has been conducted on a regular basis to incorporate a consistent and uniform "ensemble method" of adjusting correlator clock offsets at all IVS correlators to improve the accuracy of UT1-UTC estimates using VLBI. Much progress has been made in having the three major correlators (i.e., MIT, USNO, and MPIfR) use these methods. These teleconferences have tangentially improved VGOS clock-setting methods through the exposure of inconsistencies in the method of application of sampler delays.

2.9 DiFX Software Support

Support for the community continues for difx2mark4, fourfit, and HOPS. This support includes addition of features requested by users, other enhancements, and bug fixes.

3 e-VLBI

Non-real-time electronic transfers of VLBI data have continued during this period. Data from 11 stations from 20 VLBI sessions (10 R&D sessions from each of the IVS and VLBA networks) were transferred to MIT Haystack during the past two years for in-house correlation or for conversion to Mark5 media prior to shipping for off-site correlation.

4 Experiments Correlated

A total of 96 geodetic VLBI sessions have been processed, at least in part. These include 18 R&D, eight T2, 42 VGOS Trial (VT), five CONT17, two mini-CONT, and 21 VGOS-related sessions that were either broadband, mixed mode, or tests of other types. This count does not include smaller tests because they were too small to warrant individual experiment numbers.

5 Existing Correlator Capabilities

The "old" geodetic DiFX cluster, consisting of six PCs (each with a dual hex core 2.66 GHz Intel Xeon processor), was merged and integrated with the existing 16 EHT cluster PCs (each with single deca core 2.8 GHz Intel Xeon processors controlled by two equivalent master nodes) and the 38 newly purchased PCs (each with dual deca core Intel Xeon CPUs). The merged monolithic (but easily subdivided) "super cluster" has over 1,152 cores.

Connecting, providing data to, and supporting this computing infrastructure are a Gb Ethernet fabric with three 100 Gb Ethernet switches, 197 TB of new data storage space, and three file storage servers that can also act as DiFX compute nodes providing $>$200 TB of file storage. The 40 Gb/sec (Infiniband) network fabric using a Qlogic switch and 40 Gb Ethernet network fabric from the EHT cluster mentioned in the last report were retired in this process. A total of 11 Mark6 playback units with DiFX fully installed are connected to the Gigabit Ethernet fabric. In addition, new racks and uninterruptible power supplies condition and stabilize power provided from a new 208-volt power distribution fabric.

The integrated cluster (Figure 1) is used to correlate data from EHT, VGOS, and legacy S/X observing sessions.

6 Staff

The following staff have participated in various aspects of correlation, post-processing Mark6, and e-VLBI development and operations.

Fig. 1 (Left) Cluster correlator and (right) Mark6 playback units at the MIT Haystack Observatory.

6.1 Correlator Software Development

- John Barrett - software development and support
- Roger Cappallo - HOPS software development
- Geoff Crew - DiFX, HOPS, and Mark6
- Kevin Dudevoir - Mark6, e-VLBI, computer system support and development, and DiFX
- Tim Morin - cluster IT, hardware/software support
- Jon Rose - cluster IT, hardware/software support
- Jason SooHoo - cluster IT, Mark6, and e-VLBI
- Chester Ruszczyk - Mark6 and e-VLBI

6.2 Correlator Operations

- Peter Bolis - correlator maintenance
- Alex Burns - Mark6 and general technical support
- Brian Corey - correlation oversight, station evaluation, and technique development
- Glenn Millson - correlator operations
- Arthur Niell - technique development
- Don Sousa - media management and shipping
- Mike Titus - correlator operations setup, oversight, and hardware/software testing
- Ken Wilson - correlator and Mark6 maintenance

7 Conclusions and Outlook

Building, commissioning, and expansion of the VGOS network will continue. Standardization and automation of procedures to export VGOS processing to other correlators will proceed. Mixed-mode observation and correlation methods will be developed and exported. Routine geodetic processing will continue.

Acknowledgements

In memory of Kevin Dudevoir, for his contributions to the MIT Haystack VLBI correlator.

References

[*Barrett et al.,* 2019] Barrett, J., R. Cappallo, B. Corey, G. Crew, P. Elosegui, A. Niell, C. Ruszczyk, and M. Titus (2019), VGOS Data Processing Manual, *MIT Haystack Observatory*, *version 1.0*, pp. 32.

[*Deller et al.,* 2011] Deller, A. T., W. F. Brisken, C. J. Phillips, J. Morgan, W. Alef, R. Cappallo, E. Middelberg, J. Romney, H. Rottmann, S. J. Tingay, and R. Wayth (2011), DiFX-2: A more flexible, efficient, robust, and powerful software correlator, *Pubs. Astron. Soc. Pacific*, *123* (901), 275.

[*Merkowitz et al.,* 2018] S. M. Merkowitz, S. Bolotin et al. (2018), Modernizing and expanding the NASA Space Geodesy Network to meet future geodetic requirements, *J. Geodesy*, https://doi.org/10.1007/s00190-018-1204-5.