

# Mark 6 16-Gbps Next-Generation VLBI Data System

Alan R. Whitney, Roger J. Cappallo, Chester A. Ruszczyk, Jason SooHoo, Geoffrey B. Crew

**Abstract** The Mark 6 VLBI data system has been developed as a next-generation disk-based VLBI data system capable of supporting the goals of VLBI2010 and other very-high-data-rate VLBI applications, with a maximum sustained recording rate of 16 Gbps. Based on COTS data hardware and open-source software, the Mark 6 is designed to transition easily from the widely used Mark 5 system. Its features include a ‘scatter/gather’ gather algorithm to ensure that data recording is not slowed by one or more slow or bad disks. The first field demonstration of a 16 Gbps/station VLBI experiment using Mark 6 in 2012 is reported. Existing Mark 5 systems are upgradeable to Mark 6, and existing Mark 5 SATA modules are upgradeable for compatibility with Mark 6.

**Keywords** Mark 6, data recorder, broadband recorder

## 1 Introduction

In this paper we report on the status of the 16-Gbps Mark 6 VLBI data system and its readiness to support VGOS operations that are expected to start in the near future. In particular, we describe a single-baseline dual-polarization VLBI experiment with a 4-GHz aggregate-bandwidth in which data were recorded on the Mark 6 system at 16 Gbps at each station.

MIT Haystack Observatory

## 2 Mark 6 Recorder

In contrast to the Mark 5 VLBI data system, which is based on proprietary hardware components, the 16-Gbps Mark 6 system utilizes only commercial-off-the-shelf (COTS) technology. The high data rate is achieved by adopting modern high-performance motherboards and components writing simultaneously to 32 conventional magnetic disks.

The open-source software operates under a Debian Linux distribution with application software written primarily in C/C++ for the ‘data-plane’ software and in Python for the ‘control-plane’ software.

The Mark 6 (Figure 1) can accept data into four 10GigE data ports, each operating independently at up to  $\sim 7$  Gbps per port, with a maximum aggregate rate of  $\sim 16$  Gbps. The system supports SATA-interface 8-disk modules that are inserted into the Mark 6 chassis and connected to the disk controllers via COTS external SATA cables. Each external SATA cable supports four disks, so that each module requires the connection of two such cables. Depending on the recording rate, different numbers of simultaneously-operating disk modules are required. A single 8-disk module will support 4 Gbps of continuous recording<sup>1</sup>; two modules (16 disks) are required for 8 Gbps, and four modules (32 disks) are required for 16 Gbps.

A ‘scatter/gather’ algorithm for writing data to disks has been implemented in Mark 6 to ensure that the aggregate real-time recording rate is sustained even when ‘slow’ disks are present; standard RAID

---

<sup>1</sup> Recent testing of Mark 6 using disks introduced in 2013/2014 has shown that eight such high-performance disks can support sustained 8 Gbps recording.

technology is not able to guarantee sustained recording rate in the face of even a single ‘slow’ disk.

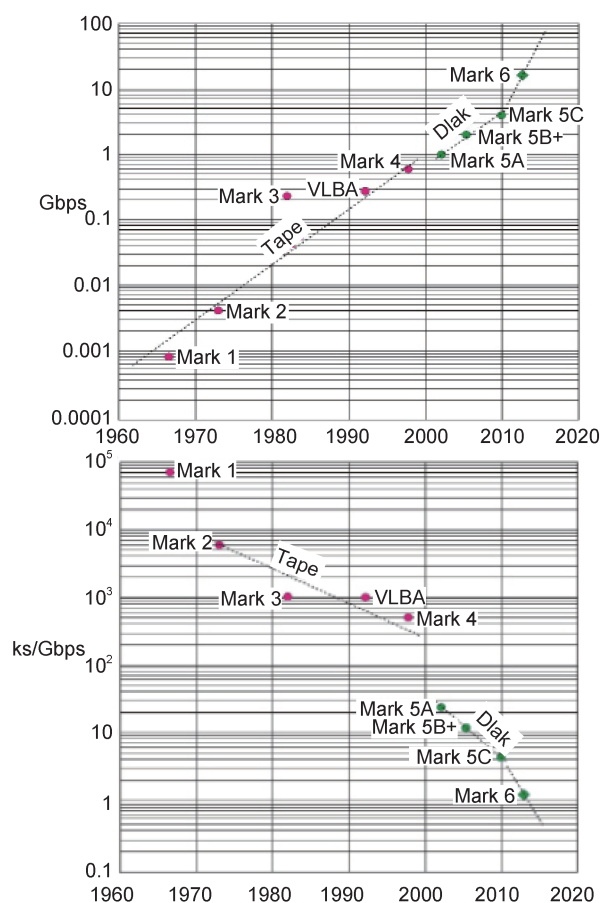


**Fig. 1** Photograph of prototype Mark 6 system.

VLBI has always pushed the technology of digital recording to the highest possible data rates primarily due to the fact that, for the majority of VLBI observations, the achievable signal-to-noise ratio of a given system increases as the square root of the recording bandwidth. Figure 2a shows the evolution of VLBI recording capability from the origins of VLBI in the late 1960s through the present. Over this period the record data-rate capability has increased by more than four orders of magnitude from less than 1 Mbps in 1967 (12" open-reel magnetic tape) to 16 Gbps today (magnetic disks). At the same time, as shown in Figure 2b, the cost per Gbps of capability has dropped by almost five orders of magnitude and moved from highly proprietary tape-based systems to semi-proprietary disk-based systems (Mark 5 series), and now to the current 16-Gbps Mark 6 system which utilizes fully COTS data hardware with specially developed open-source software.

### 3 A 16-Gbps Experiment using Mark 6

VLBI observations were made in May 2012 using the 12-m diameter antenna recently installed at the Goddard Geophysical and Astronomical Observatory




**Fig. 2** a) Evolution of VLBI recording-rate capability from 1967 to 2012, progressing more than four orders-of-magnitude from original 1967 magnetic tape to modern magnetic disks, b) parallel evolution of cost of recording in k\$ per Gbps, which has dropped by almost five orders-of-magnitude during the same period.

(GGAO), Goddard Space Flight Center, Maryland, and the 18-m diameter Westford antenna at Haystack Observatory, Massachusetts.

### 3.1 Receiver and Data System Block Design

The signal chain for each of these antennas (Figure 3) consists of a) a broadband dual-linear polarization feed; b) a broadband Low Noise Amplifier (LNA) for each polarization; c) four-way splitters for each polarization; d) four frequency converters for translation of



**Fig. 3** Diagram of 2-GHz bandwidth dual-polarization system (4-GHz aggregate bandwidth) which records to the Mark 6 data recorder at 16 Gbps.

the signals from radiofrequency (RF) to intermediate frequency bands; e) two digital backend units [2], and f) a Mark 6 16-Gbps data recorder.

Each DBE accepts the two polarizations from two bands, quantizes all four 512-MHz bands to two-bit samples, time-tags and formats the data into VLBI Data Interchange Format (VDIF) format [3], and outputs the formatted data to two separate 10-Gigabit Ethernet datastreams at 4 Gbps each, one for each band. The four Ethernet 4-Gbps streams (two from each DBE) are fed to the Mark 6 recorder through two dual-port 10-Gigabit Ethernet network interface cards.


### 3.2 Correlation and Fringe Fitting

The data were cross-correlated on a ‘DiFX’ software-based correlation system [1], processing horizontal-to-horizontal and vertical-to-vertical polarizations between the two stations. Four passes through the correlator, one for each IF band, were required to process the complete dataset.

### 3.3 Results


Figures 4 and 5 show the results of a 16 Gbps/station observation of 3C84 of 10 seconds duration. Although

there was no ability to do careful calibration of the system, the correlation amplitude and signal-to-noise-ratio of the scan (Figure 4) are within the expected range. The time-segmented band-by-band amplitude and residual-phase data (Figure 5) also appear nominal.




**Fig. 4** Correlation amplitude vs. residual multi-band delay (larger, blue, three lobed curve, scale at top) from 10 second observation of 3C84 at 16 Gbps/station on 19 June 2012; smaller, red, multi-lobed curve shows correlation amplitude as a function of residual delay rate (nanoseconds/sec, scale at bottom). Correlation amplitude is  $\sim 5.5 \times 10^{-3}$  (units in plot are  $10^{-4}$ ) with a signal-to-noise ratio of  $\sim 940$ ; HH and VV cross-correlations for all four 512-MHz bands were combined coherently for this.

Although the results of the May 2012 observations were somewhat compromised by the failure of one of the up-down converter (UDC) units, 16 Gbps were recorded for a full 60 seconds on the Mark 6 system at each station, though only 12 Gbps of the recorded 16 Gbps could be processed to obtain fringes on the weak ( $\sim 0.2$  Jy) source 0550 + 356 (Figures 6 and 7);




**Fig. 5** Plot of correlation amplitude (blue/lower dotted line) and residual phase (red/upper dotted line) vs. time for each of the four 512-MHz-bandwidth bands ('a' through 'd') and vector sum ('All') over the 10-second duration of the observation.

as in the 3C84 observation above, only HH and VV correlations were done. Manual band-to-band adjustment phases for this source were obtained from observations of the nearby brighter source 0552 + 398 (7 Jy at 5 GHz).



**Fig. 6** Results from 60-second observation of 0550 + 356 on 18 May 2012 showing correlation amplitude as function of residual delay rate in nanoseconds/sec (scale at bottom). Correlation amplitude at peak is  $\sim 4.0 \times 10^{-5}$  with a signal-to-noise ratio of  $\sim 14$ ; HH and VV cross-correlation for the three good bands were combined coherently for this result.



**Fig. 7** Plot of correlation amplitude (blue) and residual phase (red) vs. time for each of the three 512-MHz-bandwidth bands ('a', 'c', 'd') and vector sum ('All') over the 60-second duration of the observation.

## 4 Summary

The Mark 6 VLBI data system is now a mature system with approximately 30 units currently in use around the world as of this writing; the open-source nature of the Mark 6 software invites anyone to augment and/or modify its capabilities. For full documentation of the Mark 6 system, please see <http://www.haystack.mit.edu/tech/vlbi/mark6/index.html>

The Mark 6 system is available from Conduant Corporation of Longmont, CO. The cost for a full 16 Gbps (without disk modules) is  $\sim$  \$US14,000.

## Acknowledgements

The Mark 6 VLBI data system is a joint collaborative development effort between MIT Haystack Observatory, NASA/GSFC High-End Network Computing Group, and Conduant Corporation.

## References

1. 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.: DiFX-2: A More Flexible, Efficient, Robust, and Powerful Software Correlator, Publications of the Astronomical Society of the Pacific, 123, No. 901, pp. 275–287, 2011.
2. Niell, A. et al.: RDBE Development and Progress, in International VLBI Service for Geodesy and Astrometry 2010 General Meeting Proceedings, NASA/CP-2010-215864, D. Behrend and K. D. Baver (eds.), 396–399, 2010.
3. Whitney, A., Kettenis, M., Phillips, C., and Sekido, M.: VLBI Data Interchange Format — an Overview, in International VLBI Service for Geodesy and Astrometry Annual Report 2009, NASA/TP-2010-215860, pp. 15–35. (Also available at <http://vlbi.org/vdif/>)