

Progress in Technology Development
and the Next Generation VLBI System

The Mark 5C VLBI Data System

Alan Whitney

MIT Haystack Observatory, USA

Abstract. The Mark 5C VLBI data system is being developed as the next-generation Mark 5 system built on the same basic platform as the Mark 5A, Mark 5B and Mark 5B+. The Mark 5C will have the following characteristics:

- dedicated 10 Gigabit Ethernet data-input port,
- recording rates to 4096 Mbps using two Mark 5 modules (16 disks),
- full compatibility with existing Mark 5 modules,
- playback will be through a high-performance motherboard and 10 GigE NIC,
- targeted for use on software-correlator systems.

Data sources for the Mark 5C system will be based on new digital backends now being developed, specifically the DBE2/VDBE in the U.S. and the dBBC in Europe. Mark 5C data formatting is organized as one-frequency-channel per Ethernet packet, which supports an arbitrary number of frequency channels and allows fine increments in tuning aggregate data rates for better matching with available e-VLBI network bandwidths. A data mode which writes a disk module in a format compatible with playback on a Mark 5B system is planned. The Mark 5C system specification and development is being supported by Haystack Observatory, NRAO and Conduant Corporation. Prototype Mark 5C systems are expected in mid-2008.

1. Introduction

The Mark 5C is being designed as the next-generation Mark 5 system, with a capability of recording sustained data rates to 4096 Mbps. It will use the same disk modules as the Mark 5A and Mark 5B, thus preserving existing investments in disk modules.

The data interface for both recording and playback will be 10 Gigabit Ethernet, which is rapidly becoming a widely supported standard. The use of 10 GigE interfaces comes with some significant implications, however. Firstly, data sources must be designed to provide data streams in a format compatible with the Mark 5C requirements. And secondly, data playback through

a 10 GigE interface is a good match for a rising generation of software correlators. In the interests of backwards compatibility, the Mark 5C will also support a mode which writes disk modules in Mark 5B data format which can be correlated on existing Mark 4 correlators that support Mark 5B.

The Mark 5C will be implemented using the existing Amazon StreamStor disk-interface card (used in the Mark 5B+) from Conduant Corp along with a new 10 GigE-specific interface daughterboard being designed by Conduant. Unlike the Mark 5A and Mark 5B, no separate specialized I/O card will be necessary in the Mark 5C.

2. Data Sources

One major implication of the Mark 5C model is that the data source is responsible for all data time-tagging, formatting and creation of Ethernet packets. This is a departure from the VSI-H model used by the Mark 5B, which has basically only 32 parallel sample bit-streams, a clock, and 1pps tick flowing between the data source and the Mark 5B, with the Mark 5B being responsible for creating data frames with higher level time-tagging and formatting. Fortunately, VLBI data sources capable of creating such formatted Ethernet packets are now being developed in both the U.S. and Europe as part of the development of digital downconverters and backends. Suitable 10 GigE data sources to drive the Mark 5C are expected to be available sometime in the first half of 2008. Normally, each Ethernet packet from the data source will contain sample data from only a single frequency channel, although a Mark 5B-compatible data mode is specified which will write disk modules in a format that can be re-played on a Mark 5B playback unit; this will provide the ability to process the recorded data on existing Mark 4 hardware correlators.

3. Correlation

A major shift is currently developing to move from hardware-based correlators to software-based correlators, some of which already exist. Unlike the Mark 5A and Mark 5B, the Mark 5C will have no streaming hardware playback interface. Instead, the data files will appear to the user as standard Linux files and will be read as such. We expect that the standard interface for playback to a correlator will be through a standard 10 GigE interface implemented on a commercial NIC. Unlike existing hardware correlators, software correlators do not demand constant-rate streaming inputs. As such, the Mark 5C playback is well suited for interfacing to software correlators, but not well suited or intended to interface to hardware correlators.

4. General Mark 5C Characteristics

The Mark 5C will have the following characteristics:

- Mark 5C will be fully compatible with all existing Mark 5 disk modules, however some modules with older disks may limit record/playback data rates;
- at data rates above about 2 Gbps, it will be necessary to record to two 8-disk modules simultaneously, in so-called non-bank mode, which is not normally used by Mark 5A or Mark 5B/5B+;
- a 10 GigE interface for receive only will be implemented on the Amazon StreamStor disk-interface card (currently used in the Mark 5B+) by replacing the FPDP I/O daughterboard on the Amazon card with a newly designed 10 GigE daughterboard. This 10 GigE interface will be receive-only, optimized for sustained real-time recording of at least 4096 Mbps from a data source. Received Ethernet packets can be OSI Layer 2 or higher, but will only be processed by the Mark 5C at the Layer 2 level. Jumbo Ethernet packets up to 9000 bytes will be supported. The data source is required only to transmit Ethernet packets, and is not required to process any received packets;
- the entire data payload from each arriving Ethernet packet, sans a specified length of payload header (which may contain higher OSI Layer parameters or other information), will be recorded to disk. In this sense, the Mark 5C is entirely ‘formatless’; i.e. all data formatting must be done by the data source. This allows each user to format the recorded data according to his/her needs;
- the Ethernet data payload may contain a user-generated 32-bit “Packet Sequence Number” (PSN), whose position within the data payload can be specified to the Mark 5C. The Mark 5C can be commanded to a “PSN monitor mode” that will parse this serial number from every packet to identify missing or out-of-order packets. Out-of-order packets, within some reasonable limits, will be restored to proper order, while the user data from each missing packet will be replaced by user-specified “fill-pattern” data. The MSB of the PSN may also be used as an ‘invalid’ marker to prevent recording data from a packet. If “PSN monitor mode” is disabled, data are recorded to disk in the order that packets are received; no checks are made for out-or-order or missing packets;
- similar to the Mark 5A/B, the Mark 5C will record data as “scans”, where a scan is defined as the period between starting and ending the recording of a particular observation. The duration of a scan may be from several seconds to many minutes. The host application software will maintain a directory of scans for easy identification and access. No duplicate-named scans are allowed;

- scans will appear as normal Linux files to the host PC. Data playback on the Mark 5C will be through a 10 GigE NIC interface on the host PC . A planned upgrade by Conduant of the Amazon card, which interfaces to the PCI-X bus, will support the PCI e bus to allow substantially higher playback rates from Mark 5 disk modules.

5. Physical Connections

Because the interconnection between the data source and the Mark 5C is a standard 10 Gigabit Ethernet connection, there is considerable flexibility in connection topology between data-source units and Mark 5C units. Fig. 1 illustrates the general case where multiple data-source units provide data to multiple Mark 5C units through a standard Ethernet switch. The Ethernet switch allows Ethernet packets from either of the two DBE2 units to be arbitrarily routed to any of four Mark 5C units, providing an easy way to manage route data packets in an arbitrary manner, but is particularly useful to manage data-rate mismatches between individual data sources and individual Mark 5Cs. For example, an 8 Gbps packet stream from a single DBE2 could be separated into two 4 Gbps packet streams to two Mark 5C units.

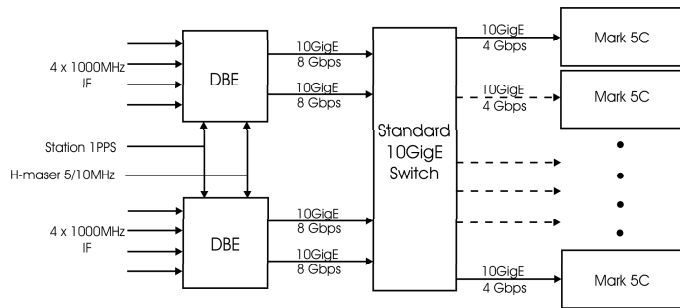


Figure 1. Generic signal-connection diagram for Mark 5C

6. Data-Organization Philosophy

Previous generations of VLBI data recorders, such as Mark 3/3A/4/5A/5B have generally formatted multiple observing frequency channels as multiple-parallel bit-streams representing simultaneous samples from all channels. Also, typically, the per-channel data rate is constrained to 2^n Mbps and the number of channels is also usually constrained to be integer powers of 2, thereby also constraining total data rates to be 2^n Mbps. This type of data organization has been satisfactory for the type of hardware correlators that have been predominant over the past 30 years, but is not well matched to today's needs. Modern VLBI data management relies increasingly heavily on the use of standardized

local and global computer networks whose capacities are not well matched to the 2^n Mbps paradigm of established VLBI data systems. Furthermore, VLBI correlation processing is moving rapidly towards fully software solutions operating on clustered computer networks. Both of these trends argue for improved flexibility in aggregate VLBI data rates and repackaging of data into more convenient forms for processing by standard computer hardware.

7. Slow-Disk Management

Considerable experience has shown a low, but significant, occurrences of “slow” disks, that is disks that still function to read and write, but at a reduced data rate due presumably to marginal or bad sectors on the disk. Normal RAID systems will simply slow down to accommodate such disks in a RAID array since they are not designed for real-time operation. The StreamStor disk-interface cards in the Mark 5 units are designed to re-distribute the load to normal-performing disks in the face of one or more “slow” disks in the recording array so as not to lose critical real-time data. Playback through a software correlator is generally non-critical as to playback rate, though the correlation speed may be slower than normal. If correlation speed is important, a playback mode may be invoked whereby data not received in time from slow disks will be replaced with a specified “fill pattern”.

8. Upgrade Path from Mark 5A, Mark 5B or Mark 5B+

The Mark 5C requires an Amazon StreamStor card (\sim \$9500) plus a 10 Gbps Ethernet daughter board (\sim \$3000–4000); no VLBI-specific I/O board is required. Upgrade from Mark 5A or Mark 5B requires an Amazon board and 10 GigE daughterboard. Upgrade from a Mark 5B+ requires only a 10 GigE daughterboard.

9. Summary

The Mark 5C development is being support by Haystack Observatory, NRAO and Conduant Corporation. Prototype Mark 5C systems are expected in mid-2008.

Detailed Mark 5C and Mark 5C data-format specifications and the corresponding data-format specifications are available at <http://www.haystack.edu/tech/vlbi/mark5/mark5/memo.html> (particularly memos 57 and 58). Other information related to Mark 5C is also available at this site.