

# Mark 5 Newsletter

MIT Haystack Observatory

December 2003

Issue #4

The Mark 5 Newsletter is issued from time to time to keep users informed regarding features, plans, problems, solutions and workarounds. All back issues of the Mark 5 Newsletter are available at the Mark 5 web site at [www.haystack.edu/mark5](http://www.haystack.edu/mark5). We invite input from anyone on subjects we should discuss or questions that need answers; please send them to [mark5@haystack.mit.edu](mailto:mark5@haystack.mit.edu).

## Contents of this Issue

1. New Mark5A software release 2.6.1
2. Important! – Make sure modules have properly written VSN's
3. Mark 5 Upgrade for NRAO "Track" Software
4. Notes on 1-Gbps recording
5. Selecting scans using the 'scan\_set' command
6. Copying scans from one disk module to another
7. Update on serial-ATA compatibility with Mark 5A
8. Mark 5B progress
9. Results of recent e-VLBI experiments
10. Limits on e-VLBI performance with Mark 5A
11. Conduant to ship modules in single-module boxes

### 1. New Mark 5A software release 2.6.1

A new updated Mark5A tar file (dated 2003 December 5 – day 339) is now available on Haystack's Mark 5 web pages (see <http://web.haystack.edu/Mark5/UpdateMark5.html>). The 'DTS\_id?' query with this version will return '2.6.1x' for command set revision. A number of bugs have been fixed and a couple of enhancements added:

#### Disk-module write-protect enhancements

During the IVS TOW meeting at Haystack Observatory in September, several people suggested that improved write protection for Mark 5 modules would be a worthwhile enhancement. Software release 2.6.1 implements an enhancement write protection system. Here is how it works:

1. A 'protect=on' command prevents any additional writing to a module.
2. A 'protect=off' command allows writing to a module to extend the current data.
3. A 'protect=off' command is required to immediately precede a 'reset=erase', 'reset=erase\_last\_scan' or 'VSN=...' command, even if protection is already 'off'. This protects a module from accidental erasure or rewriting of the VSN.

An attempt to erase a module with SSease on a protected module will elicit a query "are you sure?" for verification before proceeding. The writeup in "Mark 5 Auxiliary Programs" has been updated.

## e-VLBI protocol enhancements

UDP protocol has been added as an option to TCP for network transfers. A new 'net\_protocol' command has been added which also allows the user to specify buffer sizes:

```
net_protocol = <protocol> : [<socbuf size>] : [<workbuf size>] ;
```

See the command set at <http://web.haystack.mit.edu/mark5/software.html> for detailed information.

## **2. Important! – Make sure modules have properly written VSN's**

We have received several reports of modules arriving at correlators without an electronically-written VSN. Without an electronically-readable VSN, the module VSN information cannot be read by the Mark 5A and put into the station log. This can create problems in tracking modules from station to correlator and create delays at the correlator. Before recording any module at a station or shipping from a correlator, please check that the VSN has been properly written to it with a 'VSN?' query. If it has not been written, write it using the 'VSN=< VSN>' command.

The 'VSN' command itself has been upgraded to accept only VSN's that are in the specified 8-character format of the form 'ownerID-serial#' where ownerID is 2 to 6 characters and is strictly upper-case alphabetic (A-Z) and where the serial# is strictly numeric (0-9). The '-' indicates a parallel-ATA interface Mark 5 disk module. (Example VSN's: 'MPI-0123', 'HS-00851') The Mark5A software calculates the total module capacity (GB) and maximum data rate (Mbps) and creates the 'extended-VSN' of the form 'VSN/capacity/maxdatarate', which is written onto a permanent area of the Mark 5 disk module. (Example extended-VSN's: 'MPI-0123/960/1024', 'HS-00851/800/512')

## **3. Mark 5 Upgrade for NRAO "Track" Software** (note from Jon Romney)

“As many Mark 5 users are aware, NRAO expressed interest some time ago in upgrading our 'Track' software to handle shipments of Mark 5 disk modules. We have not been able to dedicate any personnel resources to such an upgrade until now. We are glad to report that NRAO management recently made available the 4 work-weeks which the upgrade is expected to require.

This upgrade is limited to Phase 1 of the overall Mark 5 Track capability. The primary new features to be added are tabulations of the capacity and maximum write rate for each module. Phase 2, which would add further information on the characteristics and performance histories of the individual disk drives within each module, will not be addressed at this time.”

## **4. Notes on 1 Gbps recording**

Several people have contacted us regarding difficulty in recording data at 1024 Mbps with the Mark 5A. The most usual problem has proven to be in the cabling between the Mark 4 formatter and the Mark 5A. Old cables that have been kicking around for many years and have been bent or creased do not perform well at these high data rates. If you are having trouble recording to your Mark 5A at 1024 Mbps, check that your cabling is in good shape. Make sure that all 4 cables are the same length and routed neatly from the Formatter to the Mark 5, and not strapped to any metal objects. Be careful to keep these cables separated from video monitor cables and other sources of interference. Also make sure you have terminator boards plugged into the 4 unused output connectors on the IO Panel on the back of the Mark 5 unit; if you need

terminators, contact Peter Bolis at Haystack ([pbolis@haystack.mit.edu](mailto:pbolis@haystack.mit.edu)). If there is a problem, bursts of errors can be seen on the Mark 4 Decoder while the Mark 5 is idle and configured for 1024 Mb/s recording.

Dan Smythe has written some Field System SNAP files that can be used to configure your system for 1024 Mb/s and test its performance. These “chk1024” files can be found in the ftp directory <ftp://web.haystack.edu/pub/mark5/SNAP/>, where the chk1024.\* files are for a Mark 4 rack, and the vchk1024.\* files are for a VLBA4 rack. The procedure files contain two procedures, ‘auxerr64’ and ‘trkchk64’, which are used to check the data on all tracks.

The auxerr64 procedure accumulates errors on each pair of tracks for 1 second, and reports the result. It can be used while idle, recording, or playing. If there are no errors, the response is:

```
2003.339.16:22:52.41/decode4/dqa 00000nnn 00000000 00000000
00000000 00000000 00000nnn 00000000 00000000 00000000 00000000
```

See "Mark IV Decoder Protocol" (<ftp://web.haystack.edu/pub/mark4/decoder/protocol.pdf>) for a description of the decoder 'dqa' command. The latest version of this document is in the notebook given to each participant in the recent Technical Operations Workshop.

The ‘trkchk64’ procedure can be used to check all 64 tracks of a recording. If the recording is OK, the response is:

```
2003.338.09:26:44.21/mk5/!track_check? 0 : mark4 : 64 :
2003y338d09h22m42.697s : 7332 : 0.00125s : 16.000 : n : 0 ;
```

where n is the track number.

## 5. Selecting scans using the ‘scan\_set’ command

The ‘scan\_set’ command allows you to select a particular scan, *or part of a scan*, that is subsequently honored by a ‘scan\_check’, ‘scan\_play’, ‘disk2file’ or ‘disk2net’ command. By default, ‘scan\_set’ points to the last-recorded scan. The general form of the scan\_set command is

```
scan_set = <scan name or number> : [<start play>] : [<end play>] ;
```

The optional <start play> and <end play> parameters can be specified in very flexible ways to choose any part of the scan. A few examples might be the easiest way to demonstrate the usage of the ‘scan\_set’ command:

From scan 123-1240, select a 15-second segment of data starting at 12:40:30:

```
scan_set = 123-1240 : 12h40m30s : +15s ;
```

From scan 123-1240, start 30-seconds from the beginning and scan and play for 15 seconds:

```
scan_set = 123-1240 : +30s : +15s ;
```

Play the last 30 seconds of the last-recorded scan:

```
scan_set = : -30s ;
```

A subsequent ‘scan\_play’, ‘disk2file’ or ‘disk2net’ command will play only the specified part of the specified scan. These are only a few examples; see the full description of the ‘scan\_set’ command in the Mark 5A command set memo at <http://web.haystack.mit.edu/mark5/software.html>

## 6. Copying scans from one disk module to another

A number of people have asked how to copy all or part of a scan from one module to another, either between disk banks on one Mark 5A or between Mark 5A machines. Here we will consider a few of the possibilities and give some examples:

### Copy a scan from Bank A to Bank B on a single Mark 5B

It is not possible to directly copy between the two banks on a Mark 5A unit. However, scans may be copied to an intermediary Linux disk and thence to the target module using the 'disk2file' and 'file2disk' commands, as in the following example:

bank_set=A;	Select source module
scan_set=123-1240;	Choose the scan or part of scan to be copied (see above)
disk2file=;	Creates new file with name '123-1240' in default directory
bank_set=B;	Select target module
file2disk=123-1240;	Copy scan to target Mark 5 module

### Transfer a scan from one Mark 5A to another over a network using ftp

Transferring a scan via ftp is very similar to above, except that ftp is used to transfer the source-machine Linux file (created by 'disk2file') to the target machine; at the target machine, 'file2disk' is used to transfer the file to the disk module, just as in the above example.

### Transfer scans directly from one Mark 5A to another over a network

The 'disk2net' and 'net2disk' commands are used to transfer the selected scan from one Mark 5A to another. Following is an example sending two successive scans to a Mark5A named 'target' ('T' indicates commands to target machine, 'S' indicates 'source' machine):

T:	net2disk=open;	Open socket on target machine
S:	disk2net=connect : target;	Connect to machine named 'target'
S:	scan_set=.....;	Specify scan to be transferred
S:	disk2net=on;	Transfer scan
S:	scan_set=inc;	Increment to next scan
S:	disk2net=on;	Transfer scan
S:	disk2net=disconnect;	Disconnect from target machine
T:	net2disk=close;	Close socket on target machine

David Lapsley has written a script to automate the process of transferring many scans over a network; this script and instructions on how to use it are available on the Mark 5 website at <http://web.haystack.mit.edu/mark5/downloads.html>. David has also written a script to automate the transfer of Linux files on a source machine to a disk module on a target Mark 5A, also available on the Mark 5A website.

## 7. Update on serial-ATA compatibility with Mark 5A

The new serial-ATA (SATA) disk interface is rapidly becoming available and is expected to dominate the marketplace within a couple of years. Haystack Observatory is participating with Conduant Corporation to develop the capability of the Mark 5 to support the new serial-ATA (SATA) interface disks. The goal is to offer an upgrade to the Mark 5 system that will allow interchangeability between disk modules with SATA or parallel-ATA interface (PATA) disks. We expect this new capability to be available by about summer 2004.

## **8. Mark 5B progress**

The Mark 5B VLBI data system is now being developed at MIT Haystack Observatory. It is based on the same physical platform, uses the same disk-modules as the Mark 5A, and supports the same maximum data rate of 1024 Mbps. However, the Mark 5B will incorporate a VSI standard interface and command set. This will allow VLBA systems to bypass any existing formatter and connect directly to the output of VLBA samplers (through a simple interface) at a maximum data rate of 1024 Mbps. For existing Mark 4 systems, the Mark 5B will allow connection of all 14 BBC's to two Mark 5B's for a total aggregate data rate of 1792 Mbps. In addition, the Mark 5B is being designed to support all critical functionality of the Mark 4 Station Unit, so that the Mark 5B may be played back directly to the Mark 4 correlator through a simple interface.

The VLBA correlator needs the Mark 5B system to relieve a bottleneck in the Playback Interface that currently imposes a limit of 512 Mbps in 10-station mode; with Mark 5B, the VLBA correlator will be able to process 1024 Mbps in 10-station mode. Additionally, the Mark 5B will allow the Playback Interface to be replaced with something much simpler and cheaper.

Mark 5A and Mark 5B disk recordings are not interchangeable. However, an upgrade to the Mark 5A systems is planned so that Mark 5A systems can read recordings written on a Mark 5B to create VLBA track-format output.

Prototype Mark 5B systems are expected to be available mid-to-late 2004.

## **9. Results of recent e-VLBI experiments**

Work continues to develop e-VLBI as a practical and useful data-transfer technique. Recent emphasis has been on high-speed long-distance data transfers. In collaboration with Communications Research Laboratory in Kashima, Japan, three complete geodetic experiments have been electronically transferred from Kashima to Haystack. Sessions CRF22, CRF23 and T2023 were all transferred using e-VLBI. The data was captured at Kashima in K5 format. The K5 data were then converted to VLBA format at Kashima and then transferred using BBFTP and special software to automate the transfer to a performance server at Haystack and then to a Mark5 8-pack. The main purpose of this transfer was to verify the feasibility and automation software. The data rates for the sessions were relatively low, as the path they were traversing was bandwidth limited, we expect in the future that the path bandwidth will be significantly increased as part of an upgrade. Session CRF22 involved the transfer of 442 GB of data at 11.02 Mbps using 4 parallel TCP streams. Session CRF23 involved the transfer of 489 GB of data at 11.8 Mbps using 4 parallel TCP streams. Session T2023 involved the transfer of 543 GB of data at 33.6 Mbps using 8 parallel TCP streams.

## **10. Current limits on e-VLBI performance with Mark 5A**

Significant work has been done by Kevin Dudevoir at Haystack characterizing the disc-to/from-network throughput performance of Mark5 systems. Performance is highly dependent on motherboard, kernel version, driver versions, network interface card, network subsystem configuration, and Mark5 tuning parameters (see 'net\_protocol' command). So far, the best performance achieved is 567.9 Mbps between two back to back systems with super micro motherboards (PT3DLE) using either SysKconnect 9821(copper) or 9843 (fiber) GigE NIC cards, running Linux kernel version 2.4.20 with interrupt mitigation enabled at 6000 interrupts per

second, and using disc2net and net2disc. Both Mark5's had 8-packs installed. Detailed performance studies are continuing and a report detailing the results will be published soon. For more information please contact Kevin ([kad@haystack.mit.edu](mailto:kad@haystack.mit.edu)).

### **11. Conduant to ship modules in single-module boxes**

Conduant Corporation has agreed to supply all future shipments of disk modules in *single-module shipping boxes* instead of the dual-module boxes they have been using. This should considerably help the logistics of shipping modules so that there are always enough shipping boxes available at both stations and correlators. It should also eliminate the need and cost to ship empty single-module shipping boxes to some stations to ensure an adequate supply of shipping boxes.