

Mark 5 Newsletter

MIT Haystack Observatory

June 2003

Issue #1

Now that the Mark 5A system is moving into more general usage, we think it may be useful to periodically issue a news update to keep everyone informed regarding features, plans, problems, solutions and workarounds. We also invite input from anyone on subjects we should discuss or questions that need answers; send them to mark5@haystack.mit.edu.

General Status

Currently about 30 Mark 5A units are deployed to stations and correlators; a handful of the old Mark 5P systems will soon be upgraded to Mark 5A. Conduant is now shipping Mark 5A systems; all future orders for Mark 5A systems should be made through Conduant.

Mark 5 Web Site

The Mark 5 web site at www.haystack.edu/mark5 is intended to provide a full-range of information on the Mark 5 system, including downloads for software and firmware upgrades. Please give us your feedback on this site.

Mark 5A Rev 2.5 software released

Revision 2.5 of the Mark 5A software has been recently released, which supports almost the full set of capabilities planned for the Mark 5A. Among the new capabilities of Rev 2.5:

- Important: Rev 2.5 can playback disks made with older version of Mark5A, but recordings made with Rev 2.5 cannot be played back on older versions.
- ‘Bank mode’ is now permanently on and cannot be turned off.
- All module mounting and dismounting is now handled by the keyswitches; the ‘reset=mount’ and ‘reset=dismount’ commands have been disabled.
- A new ‘protect’ command allows a module to be write-protected to guard against accidental data loss; a module must be specifically unprotected before erasure or additional writing is possible.
- A new command, ‘reset=abort’ has been added to allow disk2net, disk2file or file2disk transfers to be aborted.
- The ‘position’ command has now been upgraded so that it is active at all times, including during playback; previously it had been active only during recording or while the system was idle.
- Expanded ‘status?’ query response reports status of data transfer other than normal recording and playback to/from disk (for example, disk-to/from-network, direct to/from network, etc).
- The ‘rtime?’ query, which returns that remaining recording time on a module, now also returns the percentage of remaining unrecorded disk space on the active module.

- ‘Fill pattern’ detection is now supported on playback to allow bad or unrecovered disk data to be replaced with wrong-parity data at the Mark 5A output. This allows modules with bad or missing disks to be replayed into the correlator so that the data may be correlated with only the loss of the data from the bad or missing disks. See more details in the following article.

There are still several additional functions on which we are working:

Automatic bank-switching: This capability will allow the automatic switching from a full bank to an empty bank during recording, with the loss of a few seconds of data. There will be a corresponding capability on playback. This feature will help to eliminate the necessity to pre-schedule module changes during an experiment, so that a set of disk modules may be considered to be just a continuous set of media.

Enhanced scan directory: The current scan directory records only the scan name and length (in bytes). We are considering adding additional information to the directory, such as data mode, source name and station name. If there is any other piece of information that you consider particularly important, please let us know so that we take it under consideration.

VSN augmentation: The data written in the ‘permanent’ area of the disks where the VSN is stored will be augmented to include the serial numbers of the disks in the module. Whenever a module is mounted, this list of serial numbers will be compared against the actual VSN’s and a warning issued if a discrepancy is found.

Revised Upgrade Procedure and Host Directory Structure

A new Mark-5 upgrade procedure is now being tested. This involves a single script that uses ftp to download the tarball from Haystack, unzips and untars the tarball, reinstalls the Jungo driver, and recompiles all the Mark-5 programs with various error checks. This is intended to make upgrading easier and less prone to error.

We are also changing the Linux directory structure on Mark-5 machines to what we believe is more logical and is more consistent with standard Linux practice. Some of these changes were inspired by the organization on the VLBI Field System computers.

There will be a new login for prog (alias programmer), and this login will own all the Mark-5-related files except those that must be owned by root. A new group, rtx, comprising prog and oper, will be established, and all Mark-5-related files will be assigned to this group.

The Conduant StreamStor files will be moved from /home/streamstor to /opt/streamstor with ownership and group as noted above but otherwise unchanged. Similarly, all the Mark-5 related files now in ~jball will be moved to /opt/mark5.

Symbolic links in both ~oper and ~prog will be changed to point to the executables in both /opt/streamstor and /opt/mark5 and to the C programs in /opt/mark5. The environment variables in these logins will be changed to correspond.

A single invocation of a script file will make all these changes. Future tarball upgrades will contain this new organization, so this script will need to run before such upgrades.

Playback with Bad or Missing Disks

At the Haystack Mark 4 correlator we had now had a couple of occasions to deal with recorded disk modules with one or more missing or bad disks. Working with Conduant, we have now

developed the necessary hardware and software to deal with this sort of problem and recover as much data as possible. It works as follows:

When a pre-recorded module is first mounted for reading, information is read indicating the disks used for recording. When attempting to read, if a particular disk does not deliver its 65,528-byte data block within a specified amount of time, the data block which would have been from that disk is instead replaced with a data block containing a pre-defined 'fill pattern'. The Mark 5A I/O card recognizes this 'fill pattern' and replaces the corresponding Mark 5A data with a pattern with even parity for the duration of the fill pattern. If the correlator is configured to reject only data with the wrong parity, then just this amount of data is rejected (plus or minus a few bytes). We have recently processed an 8-disk module with one bad disk, recovering ~85% of the data, which is near the 87.5% theoretical maximum from 7 of 8 disks. Multiple disk failures within a single module will yield correspondingly less good data.

There are a couple of caveats which you need to aware of:

1. If the bad disk is a Master of a Master/Slave pair and is missing or not electrically responsive, the Slave disk of that pair cannot be accessed; this is due to the way the ATA interface specification works. In this case, the bad Master disk should be removed and the Slave partner moved to the Master position (leaving the Slave position blank); no jumper changes should be necessary since all disks should be configured (jumped) for 'Cable Select'.
2. If either the Master or Slave disk of a Master/Slave pair hangs the electrical interface inappropriately, then neither disk will be accessible. The offending disk must be identified and removed; if the offending disk is a Master, the Slave disk must be moved to the Master position.

Disk-module Conditioning

New disks fresh from a manufacturer have generally *not* been fully tested over the complete magnetic surfaces of the disk. A reserve pool of spare sectors is maintained to replace bad sectors, but bad sectors are generally not yet identified. Here is how the procedure works (according to our understanding):

1. There is no read-while-write check on disks; there are only checks on reading. Therefore, the first time the disk is written, some written sectors may be bad (but undetected).
2. If a sector is bad on reading, it is flagged as 'bad', but no other action is taken.
3. Upon the next request to write to a flagged sector, the sector will be permanently redirected to a sector from the spare pool. Depending on circumstances, the sparing process may take up to several seconds, which can create problems when the Mark 5A is operating at the highest data rates.

Therefore, in order to spare all bad sectors before a module is used to record real data, the entire disk must be written, then read to flag bad sectors, then written once more to re-direct bad sectors. Only then will bad sectors have been spared. A recommended procedure is given in the 'Mark 5A Disk-Module Assembly and Test' memo on the Mark 5 website.

Files suitable for printing these labels are available in the 'Downloads' section of the Mark 5 web site (available either in Word or pdf format).