

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

HAYSTACK OBSERVATORY

WESTFORD, MASSACHUSETTS 01886

Telephone: 978-692-4764
Fax: 781-981-0590

14 March 2002

TO: Mark 5 development group
FROM: Alan R. Whitney
SUBJECT: Revised Mark 5 disc-management procedures

Introduction

In this memo, we propose a revised disc management system (superseding Mark 5 memo 12) for the 'updated' Mark 5 system, which resides in a single chassis and uses much smaller disc carriers. The goals of the management system are, as before:

- Use software as similar to existing tape-management software as possible
- Track all discs at all times in one central database
- Allow sufficient flexibility to meet varying requirements

Labeling

Individual discs

Each individual disc carrier will carry three labels:

- 1. On the top (as carrier lies flat on a table), a label specifies serial#/capacity (in GB, where 1 GB = 10^9 bytes), as in the following example:

Top disc label [TZ072118/61] Disc serial#/capacity(GB)
barcode '-TZ072118/61' Disc model

Each disc serial number is unique and provides a natural unique identifier, analogous to the unique VSN (volume serial number) assigned to each Mark4/VLBA tape. This label is used for normal tracking and identification of each disc in normal usage. The barcode is prefaced by an identifying character '-', shown in red. The disc-serial#/capacity field has a suggested maximum of 20 characters, which should be sufficient to cover all cases.

- 2. On the bottom of the disc carrier, a label specifies the disc model, owner and date-of-purchase. For example:

Bottom disc label [Model: IBM-07N6655] Disc manufacturer-model
Owner/Date: HSTK/MAR02

It is suggested that owner names be standardized and limited to a maximum of 4 characters.

3. On the *front* of the disc carrier, visible when the disc is inserted in the Mark 5 system or in an open shipping box, a very small label indicates the manufacturer and capacity (GB). This is simply for the convenience of the user. For example:

Front disc label



Shipping box (with recorded data)

The set of 16 disk slots in a Mark 5 system is divided into two 8-disc ‘banks’ – Bank A and Bank B. In early Mark 5 prototype versions, the discs in Banks A and B operate together to support a ‘disc set’ of up to 16 discs for recording or playback. The actual number of discs in a disc set depends on data rate and capacity requirements, but will normally be 8 (Bank A filled) or 16 (both Bank A and B filled)¹. Future updates of Mark 5 will allow Bank A and Bank B to act as two ‘ping-pong’ banks of up to 8 discs each. This arrangement will support seamless recording and playback transition from bank to another.² When this capability is implemented, the standard disc set size will be 8 discs. In anticipation of this future update of the Mark 5, disc shipping boxes have been designed to hold 8 discs, corresponding to an 8-disc set recorded in one bank³.

For identification and tracking purposes, each disc set is assigned a ‘disc set ID’ at the time of recording of the form ‘SS-YYMMDD-HHMM/nn, where

‘SS’ standard 2-char station code

‘YYMMDDHHMM’ year, month, day, hour and minute of first data on disc set

nn number of discs in set (leading zero suppressed)

For example, the disc set ID ‘PT-011204-0532/8’ specifies the disc set was recorded at station ‘PT’ with the first data on the disc set taken in year 01, month 12, day 04, hour 05, minute 32 and contains 8 discs. In addition, each disc set may optionally be tagged with a 6-character experiment name or other identifier of choice. The disc-set ID and corresponding disc serial numbers of the individual discs are logged by the field system into the experiment log when a disc set is recorded. For the 16-disc sets that may exist in the early life of Mark 5, a disc set will span two boxes.


For shipping, a ‘disc-set label’ is created which contains the experiment name and barcode, disc set ID and barcode along with a listing of the individual disc ID’s (SN/capacity). For example, if there are 8 discs in the disc set, the label would be of the form:

¹ For very small experiments with low data rates, the size of a disc set could be 1, 2 or 4 discs. However, usage of disc sets smaller than 8 discs is not encouraged due to the complications of disc management.

² The data-rate capacity of an 8-disc set is sufficient to record and playback at 1 Gbps. A disc set recorded in Bank A will be replayable in either Bank A or B; similarly, a disc set recorded in Bank B will be replayable in either Bank A or B.

³ Disc sets of 1, 2 or 4 discs will also be supported, but not encouraged. In this case, multiple disc sets may occupy a single shipping box.

Recorded disc-set label
with red dot



Experiment name
(or other identifier, as desired)

Disc set ID:
Station 'PT'
Date-time
8 discs in set

Disc serial numbers


Each barcode is prefaced by a unique identifying character ('+' or '%'). One such label is placed on the shipping box plus (suggested) one extra loose label inside the shipping box for redundancy in case the outside label is damaged⁴. A red adhesive dot is applied to the label to provide an obvious visual indication of 'recorded discs', similar to manner that red adhesive dots are currently applied to recorded tapes.

For early Mark 5 usage, a 16-disc set will span two shipping boxes. Each box should be marked with an identical disc-set label along with a manual marking indicating the disc bank ('A' or 'B') in which it was recorded.⁵

Shipping box (with 'blank' discs)

When a box of 'blank' discs is shipped (usually from a correlator), it is suggested that a similar disc-set ID label be applied, for example:

'Blank' disc set label
with green dot



Target experiment name
(or other identifier, as desired)

Disc-set ID:
Shipped from 'HS'
Date, ship serial#
8 discs in set

Disc serial numbers

where 'CA152H' is the experiment name for which these discs are to be used (or other identifier as desired), 'HS' is the standard 2-char station/correlator code of the shipping origin, '011215' indicates a ship date (YYMMDD), 'XX' indicates blank discs, and '01' is a shipping serial number. A green adhesive dot is applied as an obvious visual indicator of blank discs.

⁴ If more than one disc set is contained in a shipping box (for example, two 4-disc sets), one label for each disc set must be attached to the box exterior and extra care must be taken to identify the individual discs in each disc set.

⁵ The Mark 5 system is not sensitive to the order in which discs are plugged into the playback Mark 5 system at the correlator, so an identical shipping label may be used on each of the two boxes of the disc set.

Procedures

Station

Disc-handling procedures at a station:

1. Normally, discs will arrive at stations in boxes of 8 discs each. *The station will immediately scan the two barcodes on the disc-set label on the box exterior, then open the box to scan the individual serial#/capacity barcodes (on the top label of each disc) to update the global disc database (we'll call it 'disctrack', akin to NRAO's 'track' for tracking tapes).*
2. *A set of 8 discs within a single physical box will normally stay together as a single disc set. At a station, however, if any disc is missing or broken, another disc, ideally of the same capacity, may be found and substituted (from another incomplete set, perhaps) to re-constitute a full 8-disc set for recording⁶. For maximum efficiency in disc usage, it is preferred that all the discs within a disc set be of the same or nearly same capacity; if a single disc must be substituted within a set, it is preferable that it be of same or larger capacity than the one it replaces⁷.*
3. Upon completion of recording of a disc set, it is returned to the shipping box (or boxes if more than 8 discs in the disc set), disc-set labels are printed by the field system, and the labels applied to the shipping box(es) as discussed above⁸.
4. *When a box is ready to ship, each disc-set label (on box) and corresponding disc serial#/capacity labels are scanned into 'disctrack' and the destination is specified to update the global database⁹. The box is then closed and shipped. In some cases, it may be necessary to ship a box less than full, but this is obviously inefficient and is discouraged.*

Correlator

Disc-handling procedures at a correlator:

1. Upon arrival at the correlator, *the disc-set-label barcodes (experiment name and disc-set ID) and corresponding individual disc serial#/capacity barcode labels are scanned into 'disctrack' to update the global data base.*
2. *The disc sets are entered into the local disc library, where procedures may differ from correlator to correlator. It is suggested that both the disc-set (box) label and individual serial#/capacity labels be scanned. This allows confirmation that the actual discs received agree with the station log transmitted by an independent route. The discs will most likely be re-boxed into smaller boxes for compact storage in the correlator library, but kept*

⁶ Normally, however, a set of just 6 discs is sufficient to support 1 Gbps operation, though of course at a reduced data capacity.

⁷ The data-rate capability of the disc set degrades as smaller discs are filled and the load is shared between fewer discs.

⁸ If the box contains more than one disc set, multiple disc-set shipping labels must be attached to the box exterior. The occurrence of multiple discs sets within a single shipping box is discouraged and should happen only under special circumstances.

⁹ For early Mark 5 usage of a 16-disc set, the discs from two 8-disc shipping boxes will all be associated with a single disc-set ID (identical labels on each of the two shipping boxes).

together as 8-disc sets; each box will be labeled with a label(s) similar to (or identical) to the disc-set label(s) on the shipping box.

3. When a disc set is mounted for correlation, the *correlator will read and confirm the individual disc serial numbers against the station log*. The Mark 5 system will also automatically verify consistency among the set of mounted discs.
4. When a disc set is ready for release, we suggest that both a positive indication, such as a yellow-dot sticker, be attached to the disc-set label on the library box *and* that software authorization for release also be enabled. This will help to prevent inadvertent disc release.
5. When preparing to ship, discs are removed from the library and the serial#/capacity barcode labels scanned to remove them from the library; prior release authorization must exist before the library software will allow actual release. At this point, the discs will normally be 'erased' and moved to shipping boxes, taking care to keep discs of equal capacity grouped together within each disc set.
6. When shipping, 'blank' disc-set labels are printed and placed on the box along with a green adhesive dot; the disc-set label and individual serial#/capacity labels are then scanned into the 'disctrack' database and the intended destination specified.

New-disc registration

Discs will normally be 'registered' into the 'disctrack' database when the disc labels are first created. The registration information will include disc model and serial number, capacity, owner and procurement date. Once this information is entered, all further transactions with the database require only the information on the serial#/capacity barcode label, which uniquely identifies the disc.

'disctrack'

'disctrack' does not yet exist, but the required capabilities of 'disctrack' are modeled after the highly useful Internet browser-based 'track' database created by NRAO. *The use of barcodes for data entry into 'disctrack' is assumed for data integrity*. Programmatic interaction with 'disctrack' is also possible but might impose some difficulties on stations with poor or intermittent Internet connectivity. Highly reliable barcode scanners that attach easily to a PC are available for a few hundred dollars and should be required at each station and correlator.

For each individual disc (as identified by the disc serial number), 'disctrack' will keep track of the associated experiment name and disc-set ID. In addition, if the disc has been properly 'registered' with 'disctrack', the disc owner, date-of-purchase and disc model are also tracked with each disc. 'disctrack' users will then be able to query by individual disc, experiment name, disc-set ID, station of origin, disk model, disc size, disc owner, etc.

In addition, it is suggested that some method of entering disc status along with a comment be also allowed. This could be used, for example, to tag lost or failed discs, or special conditions or instructions. Because of the large number of discs involved, tracking failed discs will provide information about disc reliability. Tracking any patterns of lost discs could help to identify the point of loss.

Barcodes

All barcodes are Code 3 of 9 with check character; this is compatible with the barcode labels currently printed on Mark4/VLBA tapes and allows the current barcode scanners to be used without reconfiguration. A single-character identifying prefix is added to each barcode, as discussed above, to provide positive identification of the barcode being scanned. The Code 3 of 9 character set includes '0-9', 'A-Z', '-.\$/+%' and <space>. Any characters in disc serial numbers or model numbers not represented in Code 3 of 9 are replaced by '-'.

Label Printing

Each correlator and station should have the capability to print any of the labels defined in this memo. The use of an inexpensive dedicated consumer label printer is suggested. As an example, we have procured a Dymo Turbo 330 LabelWriter for ~\$220 which prints a single high-quality label in about 5 seconds using a special thermal paper at a cost of ~\$0.07 per label. The labels appear to be very durable and smudgeproof. The largest standard label available for this label printer is 59mm x 102mm (2 5/16" x 4"), which is suitable for disc-set labels and is available in rolls of 300 (Dymo #30256). Disc-top and bottom labels are standard 14mm x 87mm (9/16" x 3 7/16") in rolls of 260 (Dymo #30327). Disc-front labels are very small (~10mm x 12mm or 3/8" x 1/2") and must be cut down from larger label stock.

To print disc labels, the discs can be inserted into the Mark 5 system, which reads each individual disc manufacturer, model, serial number and capacity from the disc itself. The field system will query the operator as to ownership and date-of-purchase and print the necessary labels (software not yet written).

Label Reading

In our experience, a good quality handheld laser scanning reader is by far the most preferable to read barcode labels of highly variable qualities and densities. The particular unit that Haystack uses is one of the LS 4000i Series handheld readers from Symbol Technologies, which sells for ~\$500. Less expensive units may also be satisfactory but have not been tested. Typically, such readers are available with serial, USB or 'wedge' interfaces (which go between keyboard and computer).

Label Examples

Using the Dymo label printer and label stock described above, we have created sample labels of each type described in this memo. Figures 1 and 2 show a sample disc-top and disc-bottom labels. Figure 3 shows a sample disc-set label. These labels were printed with a simple bundled Windows application program, but printing is expected to be integrated into operational software in the future. Figure 4 shows a disc carrier with labels applied. Figure 5 is a photo of the label printer used to create the labels. Figure 6 is a photo of the handheld laser-scanner barcode reader used at Haystack Observatory.

Disc Erasure

By default, when recording, the Mark 5 system always *adds* data to the end of any existing data on a disc set. A special 'erase' command must be given to reset the record position back to the origin, 'erasing' all data on the disc set; re-recording over part or all of an existing 'unerased' recording is not permitted. Normally, 'erasure' will be done at a correlator before the disc sets are sent to stations for recording. This requires that the released discs be placed in a Mark 5 system to be 'erased' before they can be shipped. It is possible that discs could be shipped 'unerased' and then 'erased' at a station just before the start of recording, but this would seem to increase the potential for accidental erasure at a station.



Figure 1: Example disc-top label



Figure 2: Example disc-front label




Figure 3: Example recorded disc-set label



Figure 4: Example of labeled disc carrier (top and bottom)



Figure 5: Dymo label printer



Figure 6: Symbol Technologies LX4000i barcode scanner