copy to
List A + B + computer group

AT/25.1.1/001

dist. 31/5/84

Data Storage and Transport Requirements for the AT

Ray Norris

23 May 1984

1.0 INTRODUCTION

The large volume of data from the AT is likely to put severe demands on the capacity of present data storage and transport technology. A discussion of this problem may seem premature because of the rapid developments in this field, but some rough idea at this stage is necessary so that related software tasks may be evaluated. By 1988 some of the arguments of this report (and certainly the prices!) will be very dated, so that this subject will need to be reviewed at intervals. However, this report makes a first stab at estimating where our priorities should lie.

2.0 THE NEED FOR DATA STORAGE AND TRANSPORT

AT data will be preprocessed at Culgoora, and then transported either to Epping or to the user's home institution for final processing. The latter option will be encouraged in order to reduce pressure on the Epping computers, but will require some sort of standard data format and medium, and may not be practicable for experiments involving large quantities of data. If the data are processed at the user's home institution, then it would be reasonable to ask that institution to provide the media or their cost: the choice of media is then partly determined by the institution, but will obviously depend on what facilities are available at Culgoora.

In either case the preprocessing at Culgoora should be such as to produce a uv data set which has been corrected for all telescope-dependent errors, and so may be processed subsequently by standard packages such as AIPS. This does not of course preclude further data correction at Epping.

A second requirement is to archive all AT data. The medium for this need not be re-usable. If it is decided to archive all data then it is logical to use the same medium for transport to Epping. Data would then be preprocessed (preferably in real time) at Culgoora, and written straight on to the archive medium. This medium is then transported to Epping, used for data processing, and finally archived.

3.0 DATA RATES

The format in which the data are encoded need not be considered here, except for the special case of a medium with intrinsically high error rates. It will be assumed that, whatever format is chosen (e.g. FITS, BACKUP), it will be encoded on the medium with an efficiency close to 100%. In the following calculation, it is assumed that each word is encoded into 2 bytes (giving 4 bytes/complex visibility), and 10s integrations are used. It is also assumed that the quantity of data other than visibilities (e.g. values of u,v,w, ha, etc.) is small compared to the quantity of visibility data. The values below are based on the correlator configuration outlined in the Correlator System Report (AT/10.4/003), and represent the total data output from the 6 telescope compact array (CA) together with the long baseline array (LBA) operating with either 4 or 5 telescopes. The maximum data rates will be increased if the planned future extensions to the correlator are implemented, but this is unlikely to affect the average data rates very much.

	LBA=4	LBA=5	
Maximum data rate	68.8	81.9	kbyte/s
Typical line data rate (512*4/baseline)	17.2	20.5	kbyte/s
Maximum continuum data rate (CA: 32*4/baseline LBA: 1*4/baseline)	0.8	0.8	kbyte/s

4.0 LONG TERM DATA RATE

The data rate averaged over a long period will depend critically on the proportions of observing time spent in different modes. One way to assess this (as the Space Telescope Institute has done) is by canvassing potential users on the types of project they are likely to propose, although even this may not reflect the proportions of projects that are finally scheduled. Here I arbitrarily assume a mean data rate of 20kbyte/s, which is likely to be an over- rather than an under-estimate. Assuming a mean 'up-time' of 240 days/year, this then gives an annual data rate of 415 Gbyte/year.

5.0 MAXIMUM DATA RATE

The maximum data rate of 81.9 kbyte/s is sufficiently low that all the media considered here will be able to record at that rate. The maximum data rate will therefore not be a problem for the recording system, although it may be a problem for the internal workings of the computer systems. The maximum data rate will be increased by a factor of about 2.5 if all the planned correlator extensions are implemented. Although all the media would still be able to record at this rate, the video cassettes are approaching their limit so that it may in practice be more convenient to record on two cassettes simultaneously. In any case, the maximum data rate does not seem to pose a problem for any of the media considered here.

6.0 AVAILABLE MEDIA

Table 1 lists the available media and their data capacities and costs. The last two rows give the cost/year for archiving all the data (i.e. assuming that each volume is used only once) and the cost/Gbyte of the media if it is used only for transportation, and can therefore be re-used an assumed 20 times. This factor of 20 is a (probably pessimistic) guess at the lifetime of a volume. The cost of the hardware (drive, interface,etc.) is not included in these final rows as it is generally small compared to several years worth of AT recording media.

6.1 Video Discs

Video discs use a laser which changes the reflective properties of a rotating metal/glass disc, and the data is subsequently read by scanning again with a laser at reduced power. Currently available discs are prohibitively expensive and can be used only once. However, projections indicate that the price is likely to fall rapidly, making them amongst the cheapest data storage media. The projected prices in Table 1 are based on a review in 'Electronic Design', and may not be reliable. Re-usable discs, in which a laser at a different wavelength can erase the disc, are under development and are apparently expected to be available soon. Current video disc drives are available that are plug-compatible with conventional Winchester disc drives. We may expect this compatibility to continue, so that using video discs offers the tremendous advantage that, in theory at least, no special software or hardware need be developed. We may in 1988 be able simply to buy a suitable system off the shelf and plug it in to our VAX.

Video discs therefore seem an extremely attractive way of storing and transporting AT data. However, their usefulness for the AT depends critically on the expected rapid fall in prices.

Table 1: Available Media

Medium	2400ft standard magtape	Video cassette tape	MkIII VLBI tape	Improved MkIII tape	Video Disc (1984)	Video Disc (proj.1988)	Erasable Disc (proj.1988)
Comments		High error rate		High error rate	Non erasable	Non erasable	Uncertain
Unit Capacity (Gbyte)	0.18	3.6	12.8	153	1.3	15?	15?
Unit Cost ($)	20	10	300	300	200	15?	23?
Cost per Gbyte ($)	111	3	23	2	153	1?	1.5?
Hardware cost (K$)	30	0.6	20+	20+	10	<10?	?
Transport cost ($/Gbyte)	5	0.2	1	0.1	150	1?	0.1?
Archive cost (K$/year)	44.0	1.8	9.2	0.8	60.0	0.4?	0.6?

The last two rows represent the cost per Gbyte of the medium if used 20 times (i.e. for transport only), and the cost per year of the medium if used for both data transport and archiving.

The standard mag tape is assumed to be 2400ft * 6250bpi * 8tracks
The video cassette is assumed to be 4hr * 2Mbit/s
The MkIII VLBI tape is assumed to be 9200ft*28tracks*33kbpi
The improved MkIII VLBI tape is assumed to be 9200ft*28tracks*33kbpi*12posns.

6.2 VLBI Media

6.2.1 Video Cassettes -

Domestic video cassette recorders have been used for VLBI (the Mk II system) for some years with great success. They have the obvious advantages of cheapness and ruggedness. A standard MkII terminal encodes digital data onto the tape, so we would not have to develop this technology, although we would have to develop an interface between the encoder and the computer. Video cassettes seem an attractive way of storing AT data because the decks are reliable and need little maintenance, and the cassettes are compact, rugged, and light, and so would ease any problems of transportation and storage.

6.2.2 MkIII VLBI Recorders -

Honeywell instrument recorders have been used on MkIII VLBI systems for a few years, and are capable of storing large volumes of data. They have a head assembly which writes 28 tracks across the width of the tape, at a recording density of 33k bits/inch. Their overwhelming advantage for VLBI, that of high bandwidth, is not relevant for AT use as all the media here can cope with the AT data rate. Their major disadvantages are that the decks require careful maintenance (more so than a standard mag tape deck) and the glass reels on which the tape is stored are heavy, fragile, and difficult to handle.

6.2.3 Improved MkIII VLBI Recorders -

Several groups, notably Alan Rogers' at Haystack, have been working to increase the recording density, by fitting the decks with narrow heads to increase the available number of tracks. The tape is then written with 28 tracks as normal, then the head assembly is shifted and another 28 tracks written, and so on. It is expected that a twelve-fold increase can be obtained this way (some claim twenty-fold), but a reliable production deck has yet to appear. If this improved deck becomes available, then it will be a cheaper way of storing data than a video cassette recorder, although whether it will be cheaper in practice (because of maintenance, storage, etc.) is not clear.

6.2.4 Error Rates -

The video cassette recorder and the 'improved' MkIII VLBI recorder both suffer from a relatively high error rate (< 0.0001) and so would need a sophisticated encoding system. However their extremely low price would still make them attractive if video discs did not show the expected drop in prices. One additional worry is that, since video cassetttes seem such a good way of storing data, why haven't commercial computer companies developed them for data storage? The reason for this may simply be that the error rates are unacceptable for normal data processing. However, for robustly encoded AT data, the loss of an occasional visibility is not likely to be important, provided the error is detected on playback.

7.0 THE CHOICES TO BE MADE

7.1 The Feasibilty Of Archiving All AT Data

The cost of archiving all the AT data will probably cost less than $1000/year, whether on video discs or a VLBI medium. A year's data (415 Gbyte) would occupy 2300 standard mag tapes, or 153 video cassettes, or 3 improved VLBI tapes, or 28 video discs (assumed to have a 15Gbyte capacity by 1988). Only the first option (which would be uneconomic anyway) represents a storage problem. It may be concluded that to archive all the AT data is realistic and feasible.

If the prices of video discs fall as projected, then they become the obvious choice for this function. Otherwise, the use of video cassettes, or possibly the improved MkIII VLBI systems, offer reasonable alternatives, although requiring sophisticated encoding to cope with the high error rates.

7.2 Choice Of Media For Transporting Data To Epping

If the data are to be archived at Epping, then the same volumes may be used to transport the data to the Epping computers. Alternatively, if separate transport media are required for any reason, the choice must be made beween re-usable video discs (if available by 1988 at a reasonable price) or the VLBI systems (video cassette or improved MkII).

7.3 Choice Of Media For Transporting Data To Other Institutions

The choice here is clearly limited by the peripherals available at the users' home institutions, and by the need to use a standard format (e.g.FITS). It is to be hoped that, if video prices fall as expected, most institutions will acquire video disc readers. Otherwise, the only alternative currently available is a standard magnetic tape. Whilst transporting a 15 day observation at the maximum data rate would require nearly 300 tapes, and so would be unrealistic, the data from typical continuum observations will fit onto a few tapes. In practice, this means that users who observe with a high data rate will be forced to process their data at Epping, unless they acquire a peripheral to read our archive format. Many users, however, will not observe at such a high data rate and will be able to take their data (or be sent their data) for processing at their home institution. In either case, it seems that compatibility with other institutions need not affect the format we choose for our archive and transport system.

8.0 CONCLUSION

At this stage we must wait and see what happens to prices of video discs. If they fall as projected, then no special action need be taken except for the purchase in 1988 of peripherals and media that can simply be plugged into our VAX. If the prices do not fall, then we must start considering the necessary software and hardware that will need to be developed in order to record data on to domestic video cassettes, or possibly onto the improved MkIII VLBI tapes.

9.0 ACKNOWLEDGEMENTS

I thank Mike Kesteven and Ron Wand for helpful comments, and Alan Rogers for details of the MkIII VLBI systems.