Proceedings of The Institute of Acoustics

DIRECT ACCESS RECORDING - A SURVEY OF THE STATE OF THE ART

F.J. Rumsey

University of Surrey, Guildford, England

INTRODUCTION

Some of the earliest digital audio editing systems used computer disk drives and solid state memory to store audio data. Recently there has been a revival of interest in the use of modern data storage media, such as the Winchester drive, for use in audio systems, and there are now a considerable number of commercial units available. In the following paper the Winchester drive as an audio storage device is examined. Following this, the limitations of Winchester storage are discussed with reference to new storage media, such as WORM and magneto-optical drives, which may form the next generation.

DISK AND TAPE STORAGE COMPARED

Disks have one major advantage over a serial medium like tape: the fact that any piece of information stored on a disk can be accessed very quickly. One only has to think of an LP record versus a cassette: in order to start playing at the required track on the record it is a quick operation to drop the stylus at the right place, whereas it is necessary to wind serially through a tape to find the start of the track, which takes longer. It is perhaps not unreasonable to suggest that this is the one over-riding factor in favour of recording audio on disks, because there are disadvantages.

WINCHESTER DISKS

The so-called "Winchester" disk drive is the type incorporated into most random access audio systems today. The Winchester drive has been used in computers for quite some time now, as it provides space for the storage of a large amount of data in a relatively small space, it is reliable, and the cost is now reasonable. The Winchester drive is different from some other hard disk drives because it is a sealed unit, and the physical disks inside it cannot be removed to make way for others. It is not like the floppy disk drive that is present on most micro-computers which allows the user to insert and remove disks at liberty.

The drive is the combination of physical disk surfaces on which the data is stored, the heads which pick up the data from the surfaces, the motor which rotates the surfaces, the servo mechanism which controls the moving parts, and the electronic section which controls the data flow to and from the surfaces and which interfaces to the rest of the system. In a floppy disk drive, all of these parts except the disk surface itself remain in the fixed system, but the storage surfaces (the floppy disks themselves) can be removed.

Figure 1

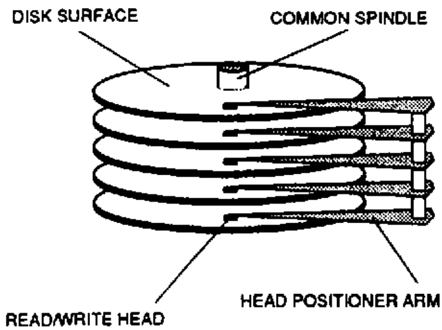

The Winchester drive must be virtually sealed (except sometimes for a small pressure relief vent) in order to prevent the surface of the disks from becoming contaminated. The lack of contamination and the fact that the disks will never be removed means that fine tolerances can be used in manufacture, allowing a larger amount of data to be stored in a given space than is possible with removable drives.

Figure 1 shows that more than one disk resides inside a Winchester drive. These are rigid, not floppy, and all rotate on a common spindle. Each surface has its own "pickup", or read/write head, which can be moved across the disk surface to access data stored in different places. The heads do not touch the surface of the disks during operation: they float just a small distance above the surface, lifted by the aero-dynamic effect of the air on the head carrier due to the rotation of the disk. A small area of the disk surface is set aside for the heads to land on when the power is turned off, and this area does not contain data.

Data and the disk

Unlike the LP record or Compact Disc, data is not stored in a continuous spiral. On the Winchester drive, it is stored in a series of concentric rings (*tracks*), each divided up into chunks called *blocks*. The term *cylinder* relates to all the tracks which reside physically in line with each other in the vertical domain on the different surfaces, and the term *sector* refers to a *block* projected onto the multiple layers of the cylinder (see Figure 2).

DIRECT ACCESS RECORDING

Figure 2

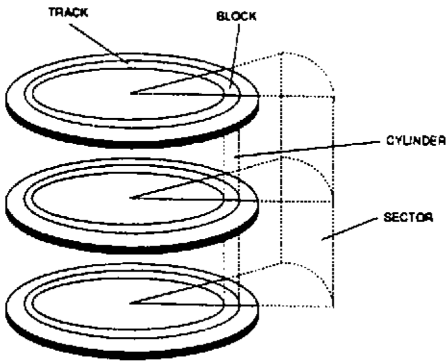

Data is stored and read in chunks which fit into these subdivisions of the storage surface, so immediately this implies that data must come off and go onto the disk in non-continuous form.

In order to "go and fetch" a particular piece of data it will be necessary to move the head on the relevant *surface* to the relevant *track*, and then to wait while the disk rotates until the start of the data section arrives under the head. So there are some delays involved:

Delay 1 - Move head across surface
Delay 2 - Wait for data to come under head
Delay 3 - Transfer data from disk to rest of system

Of these probably the longest is the time taken to move the head to the right place on the right surface. The first two delays combine (in simple terms at least) to limit what is known as the *access time*, and this can be a number of milliseconds. The speed at which data can be transferred to and from the disk is known as the *transfer rate*, and in a Winchester drive used for audio will usually be of the order of at least 8-10 Mbit/s.

Audio and the disk

One of the things that most people expect is for sound to be recorded and played back without breaks. There is an inherent contradiction between this requirement and the "burst" nature with which audio will come off a Winchester disk, because of the fact that the audio data will be spread about in different sectors requiring that there are gaps in data transfer while the heads move between sectors.

In order to produce a continuous audio output from disk a memory *buffer* is used which takes audio in bursts from disk (during replay) and which can be drawn from to supply audio samples continuously to the convertors.

Conversely, during recording the buffer takes continuous audio from the convertors and and this can be written to the disk in the right sized chunks, directed by the disk controller. The buffer is a bit like a bucket with a hole in the bottom: as long as the bucket is partly full of water it can be filled intermittently (in bursts) from a tap, yet a continuous stream of water will come out of the hole. It can probably be appreciated that the tap must fill the bucket enough during the intermittent bursts to keep it reasonably full, otherwise the water will leave the hole in the bottom faster than it is coming into the bucket from the tap; so the instantaneous flow of water from the tap must be a lot faster than the rate of flow from the hole. We could trade off a faster flow rate from the tap for longer gaps between bursts, or a slower tap for more regular bursts. (see Figure 3).

Figure 3

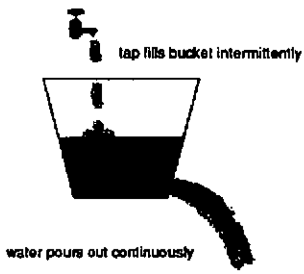

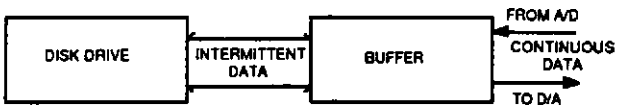

The other main point when considering the use of Winchesters for audio storage is that audio samples which were contiguous in time in the real world may not be physically close on the disk surfaces. This may at first seem quite contradictory to the requirement for fast transfer of data, because one would think that the fastest transfer of audio would occur when all the pieces of audio were located close together, as the access time delays would be minimised. Nonetheless, one must consider the situation in which pieces of audio are being edited together which were a long way apart in real time: in this case there is now a disadvantage in having contiguously recorded data blocks, especially if the disk drive is being used for more than one "real" audio track (multiple outputs, or "multi-track" operation) because the drive will be required to provide samples from a number of different locations which may be a long way apart physically on the drive surfaces. Thus the ordering of audio data blocks on the different surfaces and between the sectors is a very complex business, and not one which there is space to go into here. It is sufficient to say that a compromise must be reached between the size of audio data "units" (which may not be the same size as *blocks*), the number of contiguous samples which are recorded close together, the wastage of space on the disk surfaces, and the eventual average transfer rate that is achieved given a particular application.

One extreme is the situation in which a mono audio signal is recorded continuously for half an hour, and which will only ever be played back on its own in the same order as it was recorded, in which case there would be little disadvantage in the approach which layed down audio data blocks in physical succession on the drive surfaces. The other extreme is one in which there are a large number of small sound files stored on the disk, all of which could be requested by the system at any time and assigned to any output, and which could have any number of temporal relationships between them. In this case the idea that each file should be contiguously recorded would slow things down enormously as the pickup head would have to keep nipping backwards and forwards between the sectors that were most urgently required to feed particular outputs. Careful "scatter-storage" would be more appropriate in this case, as the physical location of sound files on the drive does not favour any one particular temporal relationship between files.

When one listens to a system in operation that uses Winchester disk drives, it is possible to hear the heads stepping backwards and forwards at high speed between cylinders as they access different blocks of data.

Audio channels and the disk drive

Another common problem encountered in the understanding of the use of disk drives for audio is that of the number of "tracks" that can be recorded. The industry is used to multitrack tape recorders which have a number of physical tracks, each of which feeds a given audio output. The disk drive is simply a large block of storage space, and no one area of this space is physically set aside for a particular track or audio channel. If a disk-based system has four audio outputs, it is a relatively easy matter to assign a sound file to any one of them by redirecting the output of the buffer to the appropriate convertor.

A disk-based system which had enough storage space to store an hour of monophonic audio could be configured to use the same space for half an hour of stereo or quarter of an hour of four-track. The limitations to this subdivision of the space into "tracks" come with the limit on the maximum rate at which the disk controller can get data onto and off the disks. One channel of sixteen bit digital audio at 48kHz generates data at a rate a little under 1 Mbit per second. It is possible to see that to handle four channels simultaneously, the drive would have to be able to transfer data at a rate of nearly 4 Mbit/s, and this increases proportionally as more simultaneous outputs are required. There will come a point at which the disk cannot output data as fast as it would be required. Going back to the bucket analogy, it is as if the hole in the bucket was now very large, and the tap couldn't fill it up fast enough to keep it full. Most current commercial systems reach this limit with around eight simultaneous outputs, and some have chosen to make the number of simultaneous outputs per drive lower than this in order to make for more flexibility in the cross-fading and editing of files. Flexibility in this area can also be achieved by the use of a greater amount of RAM buffering.

Adding more storage

Many systems which incorporate disk drives for storing audio use a common computer interface called the SCSI bus (the Small Computer Systems Interface). This means in theory that a number of SCSI drives could be added to increase the amount of storage available. The number of drives possible is often limited by the manufacturer for commercial reasons. This also means in theory that if other SCSI storage devices of a different type (say magneto-optical drives in the future) became available, these could be retro-fitted to the existing system without too much trouble. These could increase storage capacity and allow for removable disks. The SCSI bus does have some speed limitations, though, and some manufacturers have opted for other disk interfaces which can operate faster.

Many of the drives currently available in commercial systems have a capacity around 360Mbyte which will give around an hour of "disk time" which can be subdivided between channels. Often another drive or two may be added to boost this. The presentation of the "disk time" to the user is entirely up to the manufacturer, as to whether the system is configured as a multitrack dubbing system or a stereo editor.

Limitations of the Winchester

The primary limitation of the Winchester is that the stored audio cannot be removed. If the user wishes to use the system for another job he must offload the current material to a backup device, such as a tape streamer, and load up the new material. This is time consuming, and there is no way in which original recorded material can be edited directly without copying onto the Winchester. In some applications, a library of semi-permanent sound effects and music may be required as a database, to be used in a number of applications, and these might be best stored on a more permanent random access device than a Winchester, such as a CD or optical WORM (Write Once Read Many times) disk.

As already stated, the Winchester's main advantage is its speed and its storage capacity, therefore any rival removable medium would have at least to equal it in these areas. So far this has been out of the question, but optical drive technology has now reached the stage where it is being actively considered by manufacturers for incorporation into commercial systems either together with or instead of the Winchester.

OTHER STORAGE DEVICES

WORM drives

The optical WORM cartridge allows for the storage of a large amount of audio material in a write-once form. In other words, it can be recorded onto but cannot be erased. The speed of drives has not been particularly fast until now, but the second generation of drives is considerably faster and manufacturers have begun to incorporate such drives as additional storage, mainly for archival purposes. The present aim is to allow for real-time play-in of material stored on WORM during the editing process, without requiring the user to pre-load the required material from WORM onto Winchester.

The primary concern of the user regarding WORM drives is the non-erasibility, as it appears to make for greater expense in consumable costs, but the point should be made that the majority of quarter-inch analogue tape is only ever used once.

Magneto-optical drives

Within the last six months, a small number of erasable magneto-optical drives have appeared in varying states of readiness for the market. These are also being evaluated by audio manufacturers for use in random access systems.

These drives have specifications which makes them appear ideal for audio use: they are removable, they are erasable, they have a high capacity, reasonable cost and high transfer rate (as high as or higher than a Winchester). Nonetheless, there are problems to be overcome in the use of such drives. For example, the process of recording involves an erase cycle, a write cycle and a verify cycle, which makes the effective recording transfer rate much higher than with a Winchester. Even so, the likelihood of a large number of channels being recorded all at the same time is small in most current systems, whereas the greater requirement is for *output* of at least eight channels at once, which does not involve the limitation noted above.

Given the limitation on recording transfer rate, the drives would still be suitable for stereo recording usage, and this would have the advantage that the very disks used for initial sound gathering could be inserted into the editing system's drive, ready for access by the editor. This would completely bypass the stage of up- and downloading tapes to and from a Winchester before and after editing.

CONCLUSION

Magnetic Winchester drives are still the preferred medium for the random-access storage of digital audio, although they are not removable and it is not economically viable to use them as the primary recording medium. Removable optical WORM and magneto-optical drives are becoming available with attributes which make them suitable for incorporation into audio systems, and these offer the return to a more streamlined recording-editing-mastering process in which original sound material can be edited immediately without the need for copying to another medium. These devices are currently undergoing evaluation by manufacturers as there are still design hurdles to be overcome before they replace the Winchester in this field.

REFERENCES

[1] "Random Access Editing of Digital Audio";
R. Ingebretsen and T. Stockham, JAES 32,3 March 1984

[2] "Efficient editing of digital sound on disk";
C. Abbott, JAES 32,6 June 1984

[3] "The Art of Digital Audio";
J. Watkinson, Focal Press 1988

[4] "Wringing Winchester speed from erasable optical disks"
J. McCleod, Electronics, May 12 1988

[5] Digital Audio Editing 4 - Random Access Systems;
F. Rumsey, Broadcast Systems Engineering, Oct 1986