AN OPEN-ARCHITECTURE APPROACH TO SCALEABLE DATA CAPTURE

Terry Mason / Fred Thames
Avalon Electronics Ltd

ABSTRACT

The ultra high capacity disk-based data recorders now entering service offer not just a convenient and inexpensive alternative to conventional tape systems for applications like Telemetry and Flight Test but also a unique opportunity to rethink the classical models for data capture, analysis and storage. Based on 'open architecture' interface standards– typically SCSI–this new generation of products represents an entirely new approach to the way data is handled. But the techniques they employ are equally applicable to any SCSI storage device. This Paper discusses a range of practical scenarios illustrating how it is now possible to `mix-and-match' recording technologies at will–disk-array, DLT, DTF, ExaByte, JAZ, ZIP, DVD, etc.– to produce an almost infinite combination of readily scaleable plug-and-play data capture, analysis and archiving solutions. The cost and reliability benefits arising from the use of standard mass-produced storage sub-systems are also considered.

KEY WORDS

Data recording, capture, storage, disk, array, Digital Linear Tape, DLT, DTF, ExaByte, JAZ, ZIP, DVD.

INTRODUCTION

At the 1998 International Telemetering Conference, Avalon Electronics discussed how it had been possible to adapt a commercial off-the-shelf SCSI computer peripheral such as a Digital Linear Tape (DLT) drive into a flexible, general purpose data recorder for digital and analog applications[1]. Figures 1 and 2 summarise how this was done. The enabling innovation was essentially to move the problem of computer interfacing from the *reproduce* side of the classical data recording model to the *record* side of a standard SCSI data storage device (in this case a high capacity DLT drive). Once the data had been converted into SCSI format, it became very simple to manipulate and process by computer.

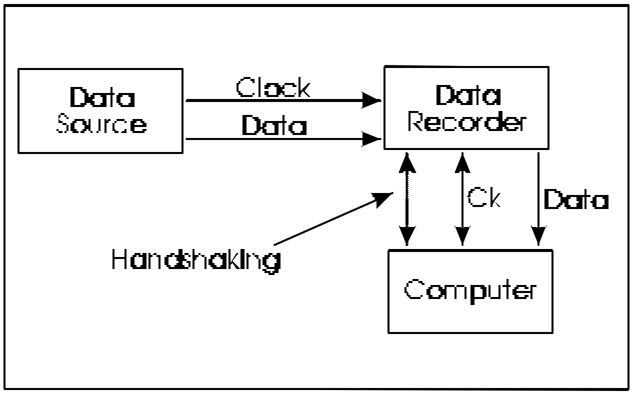

Figure 1. Classical Data Recording modes

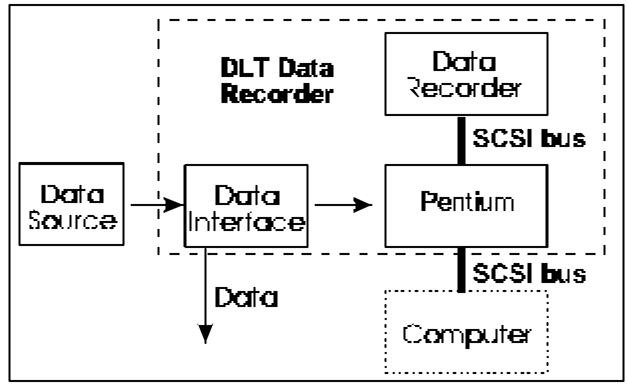

Figure 2. DLT Recorder Architecture

This concept has now evolved from its simple `stand-alone' roots into an entirely new *scaleable* approach to the way data is captured, analysed, transported and archived. Furthermore, using essentially the same schematic diagram–but with alternative, faster storage devices such as disk arrays–it has even been possible to push sustained data transfer rates beyond anything currently possible with magnetic tape. But when these extremely fast data capture devices are fitted with generic analog and/or SCSI interfaces data can readily be migrated to less sophisticated storage media for subsequent transportation, analysis or archiving. This 'horses-for-courses', scaleable architecture mean that it is no longer necessary to equip an entire data handling network with expensive storage equipment. The disk array implementation is used as the foundation for this paper in order to demonstrate the full possibilities of this approach.

HIGH RATE DATA CAPTURE USING DISK TECHNOLOGY (AE7000)

The storage 'engine' of the 1 Gbit/s Avalon AE7000 disk recorder (Figure 3) is an array of eight 18 GigaByte (GB) hard disks, giving a total storage capacity of approximately 144 GB (1.15 Terabits), or almost 50% more than the largest ID-1 tape cartridge. The unit's architecture (Figure 4), which is the subject of a number of international patent applications, is similar to that of the DLT recorder (Figure 2) in that data (either analog or digital) is first converted to 16-bit SCSI format in a Data Interface controlled by an embedded Pentium computer. In this case however, the formatted data is fanned out over the array of disks which appears to the Pentium as a single SCSI "`Direct Access" device. The inherent integrity and reliability of today's disk technology means that there is no need to operate them in *redundant-array* mode for most applications; leaving virtually the entire capacity of the media available for data storage. However, in the Avalon implementation the redundant-array mode is nevertheless retained and can be selected if required. In this case, the eight disks are divided into two groups of four, with one disk in each group acting as a simple exclusive-OR 'parity track'. The price for this additional (and generally unnecessary) security is a 25% penalty in both capacity and maximum

throughput. Input/Output interfaces are modular, the range at present including serial and parallel

Figure 3. AE7000 1 Gbit/s Disk Recorder.

digital interfaces with aggregate sustained throughputs to 1 Gbit/s, single channel wideband analog inputs to 50 MHz and flexible mixed analog and digital interfaces for telemetry and similar multi-channel applications. To overcome the twin needs to store sensitive data securely when not in use and to be able to install extra capacity when required, the disk array is housed in a single hot-swappable crate. This can be locked away in an approved safe. The various methods of reading data are discussed later.

DATA MANAGEMENT

One of the most difficult aspects of adapting computer storage devices for data capture applications is data management. Invariably, ‘really world’ data is some form of time history where the instantaneous changes of one or more physical parameters (pressure, temperature, force, amplitude, etc.) are ‘plotted’ against a timebase. Multi-track linear magnetic tape recorders are particularly suited to this application, since the timebase is effectively the tape movement itself. Helical and transverse scan recorders operate in a similar manner (albeit with some temporary manipulation of the data to accommodate head switching). Computer peripherals behave somewhat differently–typically writing in a disjointed manner which bears little or no relationship to ‘real time’. Yet many potential users, while attracted by the advantages of peripheral-based solutions, are nevertheless principally concerned in what is happening to their signals between specific points in time. This problem is compounded when edited portions of the data are transcribed from one medium to another–say from the disk array to a JAZ disk. It is invariably the *original time history* which is of interest, not the time when the data was transcribed.

In the Avalon approach (Figure 4), these difficulties are overcome in an ingenious yet simple way. The basic unit of storage on a disk is the *sector* (512 Bytes). In order to minimise the amount of overhead written to disk, input data is cached until there is enough to write a *cluster* (defined as 1,000 *sectors,* or 512 kB). As the eight disks effectively write or read in parallel, the basic unit *of system* storage is therefore 8,000 sectors (4.096 MB). Each time a cluster is written to disk a directory entry is made as an ASCII space delineated file with entries for cluster number, time (in seconds since 1980) and any related user event marks. Each time the recording process is terminated–always at the end of a complete cluster– the complete directory is written to a flash memory which is integral to (and therefore travels with) the removable disk crate. The time source can be the unit's internal generator or one of a number of external timecode options (IRIG, GPS, etc.). To avoid the complexities of multiplexing any low rate *housekeeping data* (range timecode, voice annotation, etc.) into the primary high rate data stream, these 'tracks' are stored in flash memory too, suitably linked to the timeframes to which they pertain.

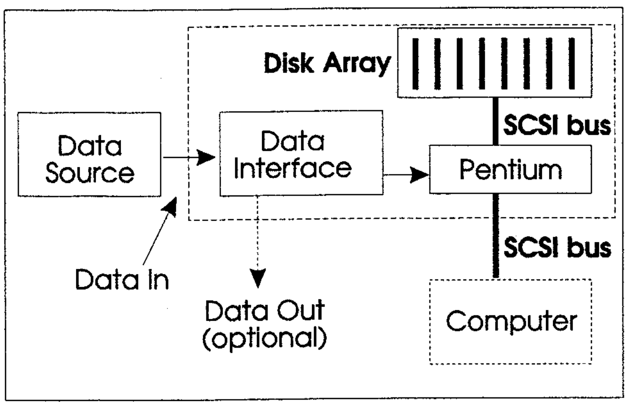

Figure 4: Disk Array Architecture

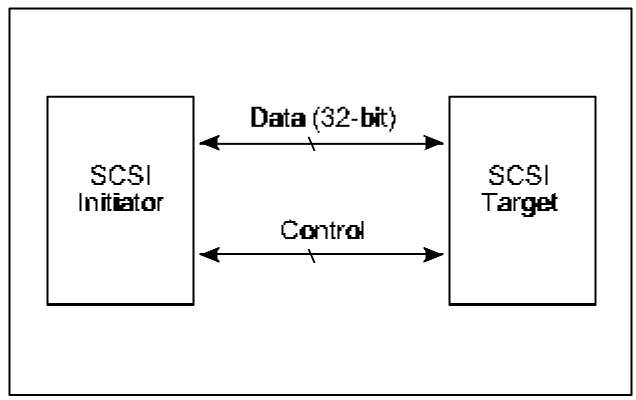

Figure 5: SCSI Data/Control Interface (fast/wide implementation)

One of the advantages of using the SCSI interface (Figure 5), rather than say Ethernet, is that it has no protocol layer, meaning that data can be simply manipulated as a variable number of blocks. All SCSI interface vendors provide drivers that bring the SCSI device to a common software interface such as ASPI (Adaptec SCSI Protocol Interface). Handshaking commands and responses (ATN, BSY, ACK, etc.) are transmitted between the Initiator and Target on separate dedicated lines. Another major advantage is the limited number of commands needed to move data around the system. For example, consider the need to transfer to RAM a passage of data of known duration which was recorded commencing at a certain moment in time. It is simply necessary to convert this *start time* into seconds (from 1980), relate this to its associated *cluster number* in the flash memory and then multiply this value by 8,000. This is the *start sector* which is then placed in the SCSI *command block* along with the appropriate transfer length (expressed in *blocks*). Once transferred, this data can be given a filename and resaved, if required.

The foregoing scenario described the use of AE7000 as a SCSI *Target* in which it appears to the host as a single 144 GB hard drive. But it is probably in its role as an *Initiator* that it makes its greatest contribution to the concept *of scalability.* As an

Initiator, it is able to download portions of data to any other suitable SCSI archival or transfer medium including, DLT, DTF, ExaByte, JAZ, DVD or other large capacity disk. Typically, in this mode the beginning and end data points are selected manually using a pair of cursors at the unit's graphical interface (Figure 6). Pressing the TRANSFER button causes the system's control computer to make the proper calculations and then transfer the dataset to the selected data in its original form. For example, an important passage of analog data could be transcribed from one AE7000 in computer compatible form onto, say, a JAZ disk for transportation to another location where it could be imported to a second AE7000 and displayed in its original analog form again. The normal error checks associated with data interchange between SCSI devices apply, meaning that no errors are introduced during the transcription process.

SYSTEM INTERFACING

Before going on to consider the range of practical implementations, it may be helpful to describe several important system interfacing possibilities. AE7000 is designed initially to operate as a stand-alone recorder/reproducer with a minimum of user controls. For many applications, this is how it will be used. Recorded data will be monitored and analysed in real time and only rarely saved for further processing later. The primary pushbutton controls are: STOP, RECORD, PLAY and GOTO, the last being used in conjunction with a pair of cursors to define the start and finish points for a passage of data to be replayed or transcribed (Figure 6).

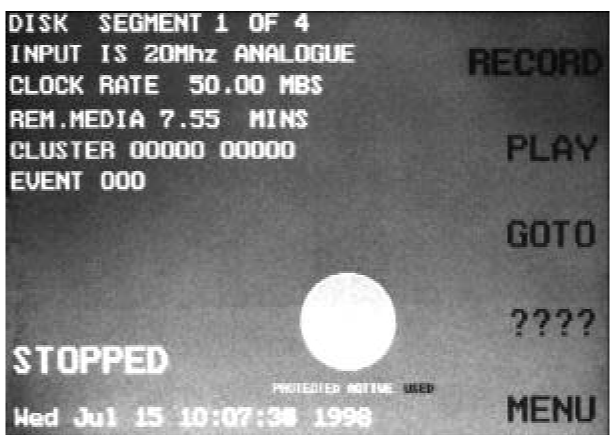

Figure 6: AE7000 Control Screen

But in the context of a scaleable data capture strategy, a number of more adventurous possibilities exist (Figure 7). Bearing in mind that disk-based systems are essentially *temporary* storage devices, it is important that there should be a convenient means of copying data for transportation, off-line analysis and archiving. As many applications are still (and are likely to remain, at least in the foreseeable future) analog in nature (telemetry, SIGINT, etc.), the simplest possibility is to transcribe selected passages in analog form to a suitable analog tape recorder. An alternative is to transcribe data directly to a standard

SCSI device. Virtually any SCSI system can be used, provided that it has the required capacity and transfer rates for the data sets involved. Since AE7000 can operate as either an Initiator or a Target, analog (or digital) data stored on a SCSI device can at any time be up-loaded again and restored to its native format.

It is an often-heard truism that 95% of all stored instrumentation data is 'garbage'. Yet any missile test range or flight test center is likely to have a massive warehouse of large, expensive, open-reel and/or cartridge tapes, only a tiny fraction of whose contents are classified as 'valuable' data. It will be immediately obvious, therefore, that the flexibility of the disk-capture approach offers important advantages in this respect. Data acquired at rates beyond the reach of even the fastest tape recorders can readily be transcribed to any suitable low cost medium provided only that the passages of interest can be identified by time of acquisition and/or simple event marker. A 5.25 inch slot is provided within the unit for this purpose, although the user is often more likely to employ higher capacity media such as DTF or DFT for larger datasets.

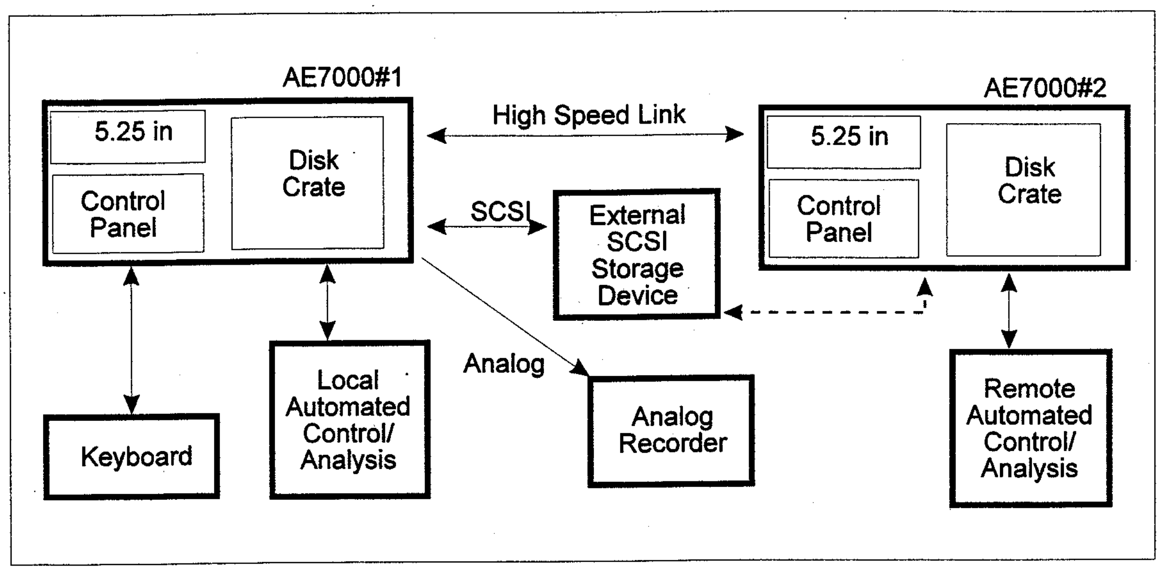

Figure 7: AE7000 Typical Data Flows

An interesting extension to this is the concept of *remote data capture*. In this case, data is captured on disk at the acquisition site while selected passages are transmitted to a remote location over a high speed link for further analysis on a similar disk recorder. This concept is particularly relevant in applications where several remote acquisition sites are controlled by a single central monitoring facility.

DISK-BASED DATA CAPTURE IN TELEMETRY

From the standpoint of the capture and analysis of telemetry data, some important advantages of a scaleable disk-based solution are:

1) ***'Instant-on'* buffered data capture**. There is no need to run the recorder up to speed in anticipation of the arrival of data. Take the case of an airborne test firing of a missile. It is normal for the aircraft to make one or more circuits of the range while the missile's instrumentation packages are calibrated, checked and their outputs recorded for future reference. Then, typically, there will be a pause before permission is given to launch the missile. There is the permanent dilemma as to whether to leave the tape(s) running (with the possibility of running out before the actual firing), or alternatively to stop the tape and hope to be up to speed again in time for the test. These problems do not exist with a disk recorder which can record instantly up to its maximum data rate (1 Gbit/s aggregate or 50 MHz aggregate in the case of AE7000). Capacity is also unlikely to be an issue since the unit's 1.15 Terabit storage is probably more than adequate to hold a full day's data at the most busy range facility.

2) ***'Loop store'* mode.** The previous comments assume a typical scenario where the data capture process conforms to a pre-planned schedule–and in any event in which personnel at the recording facility have some control over events. There are many situations however when the exact timing of an event of interest are outside the control of the recorder operator. Where high rate tape recorders are used to capture such unpredictable events, it is normal to use two or more units working in series–with someone in attendance to change the tapes as necessary. A disk-based recorder can overcome this problem by working in *loop store* mode, whereby as the medium becomes full the earliest data is overwritten ensuring that it always holds a 'loop' of 'pre-event' data. In this situation it is normal to further program the system to stop recording at a predetermined interval after the event of interest has been identified, freezing a finite time history leading up to and following the event for analysis.

3) ***Instant Access* to recorded data**. There is no 'rewind time' with disk recorders. Once a test run has been recorded it is available *immediately* for verification and/or analysis. The operator simply has to select the timeframe of interest using the cursors and then press the GOTO button. *Looping* mode is also available on replay, allowing an event to be played again and again without stress to the instrument or recording medium.

4) ***Data Distribution***. Once initial validation of the data has been confirmed it is available for immediate transcription in analog or digital form (as appropriate) to another storage device, (Figure 7). This is another area where the flexibility and scalability of

the concept is important. Provided only that a human or automatic 'operator' at the acquisition site is able to identify the data of interest, it may immediately be reduced in volume to something as manageable (and inexpensive) as a medium format cartridge or JAZ disk. Nothing is lost however, since any data migrated to a SCSI or analog device in this way can always be restored to its original form, simply by reversing the process and returning it to the disk recorder. Secondary copies can be distributed to the Test Range's Client, if required. Alternatively, the data can be transmitted over a high speed data link to a remote location for further analysis.

5) ***Cost***. Compared to 'traditional' data recorders, the relatively low cost of the disk system itself, the analog and SCSI devices it is designed to interface with and the media they use offer immediate advantages. Less obvious perhaps, are the cost savings surrounding the introduction of the package into an existing data collection facility. For example, the multi-channel data input/output interfaces can be made to emulate exactly the user's existing interface (IRIG, DCRsi or ID-1, for example). Where a possible change of recording technology is linked to a move from an analog to a digital environment, the traditional horrors of 'going digital' are eliminated. AE7000 for example, fits equally well into either regime, and will perform the analog-to-digital conversion (i.e. to SCSI) 'for free'.

6) ***Reliability and Logistics***. The principle moving parts in a disk array system are the disk drives themselves, and these are effectively field-replaceable 'throw-away' items. Other SCSI devices in the network can be considered equally expendable, or at least inexpensive enough to carry an adequate number of spare units. Furthermore, experience has shown that the inherent reliability of mass-produced computer peripherals can be retained when they are adapted for field data capture applications provided care is taken in the re-engineering process. For example, Avalon Electronics offers a 5-year warranty on the disk drives used in its AE7000 product. These considerations mean that expensive maintenance contracts and high-value spares holdings could soon be a thing of the past as the instrumentation world embraces these new technologies.

CONCLUSIONS

Over the years there have been many 'false dawns' in the quest for COTS (commercial off-the-shelf) solutions to complex data capture, manipulation and archiving problems. In many cases, tape and disk-based products have failed to deliver on their initial promise. Now, however, it can be seen that a whole family of SCSI-based devices hold the key to a new, genuinely flexible approach. Compared to traditional recording tools they are inexpensive to own, simple to integrate into existing set-ups and require little or no special software code. But above all, they offer a degree of scalability to the problem unimaginable just a few years ago. As has been shown, the concept is as applicable to a simple ZIP or JAZ drive as it is to a massive 1.15 TB disk array–indeed an identical proprietary software package could probably control either. It is predicted with some certainty therefore that multi-technology solutions, based on common generic interfaces like SCSI, represent the ‘next generation’ for many data storage applications.

---

[1] F.Thames Jnr, “A New Generation of Data Recorders based on DLT Technology”, Proceedings of International Telemetering Conference, 1998.