COMPUTER-FRIENDLY
HIGH RATE DIGITAL CASSETTE RECORDERS

Edwin A. Kayes
Penny & Giles Data Systems Ltd
The Mill, Wookey Hole
BA5 1BB, England

ABSTRACT

The world of instrumentation data recording has traditionally been concerned with recorder performance in terms of bandwidth, data rate, tape speed and recording time, with the apparently unceasing trend to record more and more data.

However, while this may remain a valid perspective for data acquisition, the increasing requirement to integrate equipment into computer based environments has resulted in the need for greater emphasis to be applied to such parameters as data control and interfacing when specifying digital data recording systems.

This paper addresses these operational issues and describes the practical implementation of a computer friendly digital cassette recorder which provides a common platform for both high rate data acquisition and computer based data analysis.

Key Words: Data recording, High-Rate Digital Cassette Recorders

INTRODUCTION

The magnetic tape recorder has played an essential role in the capture and storage of instrumentation data for more than thirty years. During this time, data recording technology has steadily progressed to meet user demands for more channels, wider bandwidths and longer recording durations. When acquisition and processing moved from analog to digital techniques, so recorder design followed suit. Milestones marking the evolution of the data recorder through these various stages - multi-track analog, high density longitudinal digital and more recently rotary digital - have often represented important breakthroughs in the handling of ever-greater quantities of data.

Throughout this period there has been a very clear line of demarcation between data storage methods in the “instrumentation world” on the one hand and the “computer peripheral world" on the other. This is despite the fact that instrumentation data, whether analog or digital at the point of acquisition, is now likely to be processed on a digital computer at some stage. Regardless of whether the processing device is a small personal computer, a work-station or the largest supercomputer, system integrators have traditionally been faced with the same basic problem - how to interface what is essentially a manually controlled, continuously running device (the tape recorder) into the fast start/stop computer environment without resorting to an excessive amount of complex custom interfacing and performance compromise.

Computer friendliness also implies reliable and convenient data management. It is relatively easy to append housekeeping data during recording, but what type of data will be most useful, and how can it be used to best effect? For example, if the user intends to search his records by date, time or event, it is critical that he develop an overall strategy for the creation, logging and management of this type of auxiliary information.

BRIDGING THE GAP

Traditional multi-track recorders (both analog and high density digital) take the timebase of the information to be recorded for granted. The tape runs continuously at an appropriate speed and data is applied to the input for the duration of the experiment or process. If a recorded tape is re-wound and replayed at the same speed and in the same direction, the output is expected to be a close representation of the original input data, including its timebase. Timebase compression or expansion can be achieved by increasing or decreasing the tape speed. Time inversion is also possible by reversing the direction of tape movement. The important point is that an indication of the passage of time is inherent in the operation of the classical data recorder.

Figure 1. Continuous Recording.

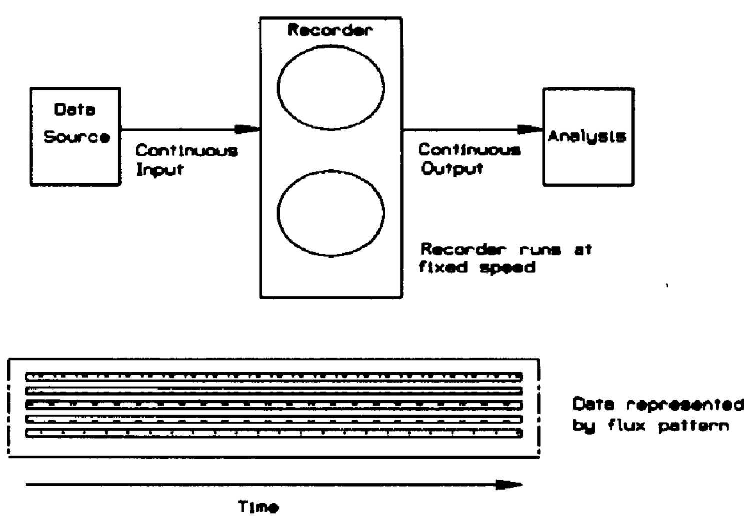

Until now, this feature has been both a strength and a weakness. A strength in terms of the ability to manipulate the passage and direction of time on a recorded experiment during the analysis process, but a weakness when it is necessary to input the data to a computer in anything but the simplest free-running mode. Given that most computers require data to be input to disk or memory in chunks at a fixed rate, it is not a simple matter to control the data flow from a constant speed system without recourse to time-consuming stop-reverse-restart routines. In contrast, computer peripherals start and stop rapidly in order to control the flow of data. This latter attribute would, therefore, appear to be a necessary characteristic for a data recorder to be considered as computer friendly.

In addition to fast start/stop of the tape itself, some high rate digital cassette recorders incorporate input and output data buffering to allow the tape transport to start and stop during data transfer as necessary. The buffer capacity will be determined by the need to ensure that all possible sequences of tape movement (ramp up, ramp down, etc.) can be accommodated without loss of data.

The use of buffered data input/output, while greatly simplifying the actual transfer of data, introduces more wide-ranging implications than might at first be obvious. For a user to gain the maximum benefit from the closer integration of the recorder into the computer environment, it becomes necessary to consider the whole data acquisition and analysis process rather than just the recorder itself.

If we accept the fundamental principle that computers need to clock data into memory in bursts by starting and stopping the tape, how are we going to retain the important timebase information which was so conveniently available by the very movement of the tape on a continuously running system? This consideration leads naturally on to the actual control of data. On the command to start, traditional data recorders ramp gently up to speed, lock in and then data is available on the correct timebase. When told to stop, they ramp gracefully down again to rest. If "good" data has been recorded on the tape at these ramping points, it is effectively lost or at least corrupted due to the stewing of the tape speed.

BUFFERED DATA TRANSFER

It is most unlikely that the clock rate of the acquisition process (e.g. analog-to-digital conversion) will be identical to that of the analysing computer. This means that a change of timebase is almost certain to be required somewhere within the data path. Looking at the complete system, several important points should be considered. In any recording system, if the tape is to be used efficiently data should be recorded on tape at the maximum rated density.

In the case of a continuously running system (longitudinal or rotary), this has traditionally meant adjusting the tape speed (and scanner speed, if appropriate) to match the input or output data rate. However, when the recorder incorporates a read/write buffer, it is usually arranged so that data is written to or read from tape at a single, fixed rate and tape speed.

Input/Output rates below the recorder's specified maximum will result in its buffer filling or emptying at a slower rate. The recorder accommodates this by automatically stopping the tape until such a time that the level of data in the buffer reaches a pre-determined level. The rate at which data is written to, and read from tape is, therefore, completely independent of user data transfer rate. This severance of the traditional direct link between user data transfer rates and tape read/write rates means that a buffered system can also accommodate data which is not continuous (i.e. intermittent or burst data) and be able to operate at any user controlled transfer rate (continuously variable) within its rated range.

Figure 2. Buffered Data Transfer.

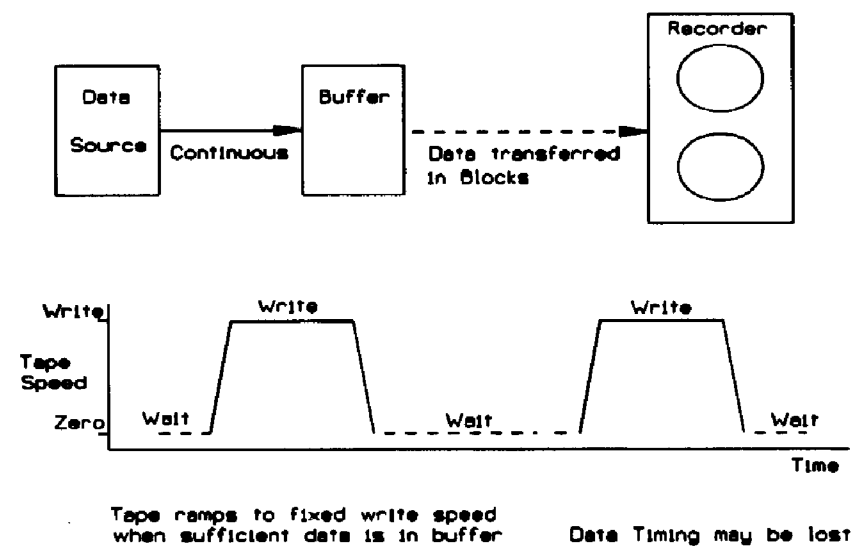

Clearly, the buffered approach would appear to have important advantages for computer based applications, particularly if the tape drive is specifically designed for very fast start/stop operation - thereby necessitating only a relatively small data buffer.

An interesting additional benefit, which should not be overlooked, is that buffered systems do not have to actually be in the normal recording mode (with tape running) in order to capture, say, an unexpected transient event. They can wait in standby mode until the event commences and then data can be written to tape from the buffer as previously described. This reduces wear and tear not only on the recorder itself, but also on heads and media in the case of fixed-head systems where nothing is in motion until data is transferred from the buffer on to tape.

Similarly, when reading data at a low transfer rate, tape motion only occurs as necessary to maintain a level of data in the buffer commensurate with the user transfer rate.

DATA FORMATTING

Intuitively, it would seem desirable to establish a common data format throughout the data capture and processing path if only to avoid the complexity and cost of unnecessary format conversions. This philosophy requires an analysis not only of the way data is to be recorded, but of the whole network (both current and planned future expansion) to establish, for example, the best word width to use (for example: 8, 16 or 32 bits). Some recorders support only 8-bit formats while others can be user configured for all three formats. If a common interface format can be used throughout, the total system can be greatly simplified.

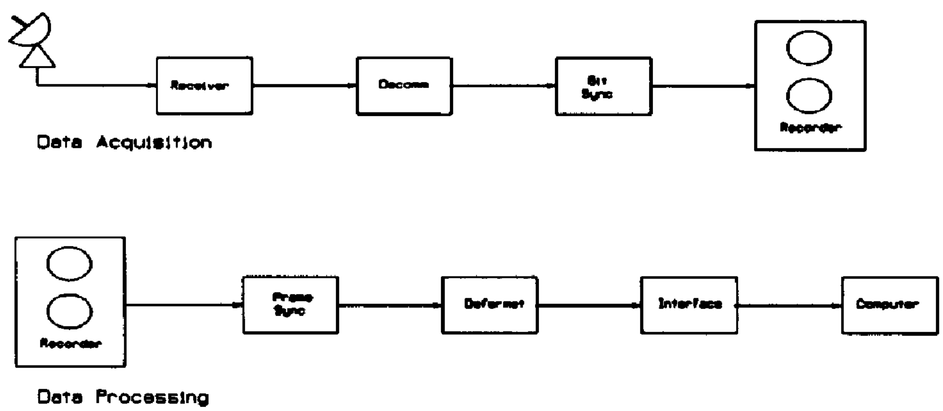

Figure 3A. Signal Path Using Traditional HDDR

If the source data is serial in nature, it is important to decide carefully when to convert from serial to parallel. In general, high rate serial recording channels are complex and expensive, so it is often best to perform the conversion before recording. The idea of standardizing on a common data interface format will generally reduce overall system complexity and cost, with the added benefit of increased flexibility and equipment utilization.

Figure 3B. "Computer Friendly" Signal Path

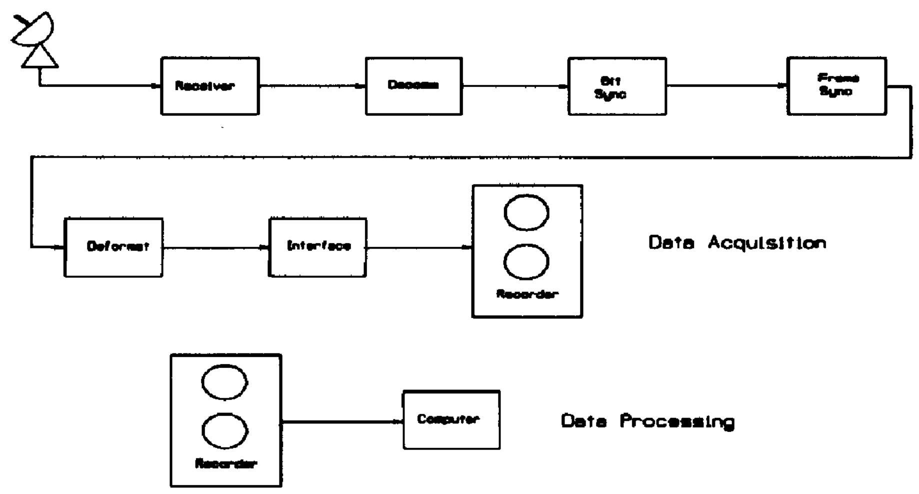

AUXILIARY DATA

While we have seen that the buffered approach has much to commend it with regard to the handling of different (and perhaps variable) input/output transfer rates and computer entry, there remains the problem of the consequent loss of relationship between timebase and tape motion since, as we have already discussed, the tape only moves when data is passing between tape and buffer. If timing is already intrinsic in the user's data stream - for example, where the input clock is synchronous with the analog-to-digital sampling process - only periodic up-dates may be necessary in order to keep everything under control.

Alternatively, more precise timing information may be required. Some high rate digital cassette recorders incorporate an internal clock which is written to a separate (auxiliary) track in the form of a date/time code. This timing information may subsequently be used to support high speed search during replay.

Another useful method of providing reference information is by using event markers. On some recorders, the controlling computer can write unique event markers along with event ID character strings to the auxiliary track. These can be scanned at high speed in order to locate selected records and also to provide an event log or directory of all events on a tape. With buffered systems, users should expect this information to be recorded in synchronizm with its associated user data in order to maintain the necessary precise relationship between the location of the event marker and the data to which it refers.

Figure 4. Auxiliary Data.

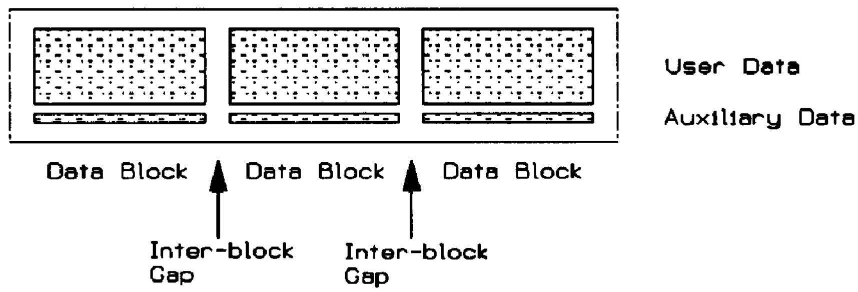

Auxiliary Data:

* Internally-generated timing.
* Timing can be referenced to an external source
* Externally-triggered Event Marker, with annotation

COMMAND AND CONTROL

Clearly, significant improvements over traditional methods of control of data recorders are needed if systems are to be integrated successfully into the computer environment. To illustrate what is now possible, command and status sets for a commercially-available recorder are shown in Tables 1 and 2. In these examples, the use of plain-language syntax greatly simplifies the preparation of control routines. Typically, commands and status requests pass between the recorder and controller via a conventional communications interface such as IEEE488 or RS-449.

DATA FLOW

The control of data flow in continuously running systems is relatively simple since it is only necessary to start the tape running (at the correct speed) and allow data to flow in or out of the recorder. With buffered systems, however, the movement of the tape itself is a secondary issue as this process is automatically controlled by the action of the recorder attempting to empty or fill its buffer. One advantage of a recorder which has been designed with an integral buffer is that it should not be possible to either overfill or empty its buffer during data transfer operations.

With continuous inputs, this may simply mean ensuring that the input clock rate does not exceed the rated maximum for the recorder. If the input is in the form of burst data - blocks of finite length with gaps in-between - it is generally permissible to exceed the maximum continuous rate for short periods. In the case of such “burst” data, it is advisable to implement a “hand-shaking” protocol so that the recorder can control the flow of data within the capacity limits of its buffer.

Table 1 - Typical Command Set

(Manny & Giles PEGASUS)

COMMAND	PURPOSE
Immediate Commands	
CLOSE DOOR	Closes cassette door after loading or un-loading.
CYCLE	Outputs current contents of Data Buffer continuously.
DATA INTERFACE n	Selects Data Interface (word width, serial, etc.)
EVENT MARKER n,	Tags write block with unique event "n" plus
user_aux_data	(optional) user Identifier.
OPEN DOOR	Opens cassette door for loading or unloading.
READ FORWARD	Starts reading from tape. Optional forms:- READ FORWARD [Start Event Marker], [Stop Event Marker] READ FORWARD [Start Time Code], [Stop Time Code]
READ REVERSE	Enables data to be read in reverse.
SET DATE	Sets internal calendar.
SET TIME	Sets internal clock.
SLEEP	Puts recorder into standby.
STOP DATA	Stops writing or reading process.
WRITE	Puts recorder in write mode.
Deferred Commands	
BOT	Returns cassette to BOT (Beginning of Tape).
EEMODE	Performs diagnostics routine.
FIND EOD	Initiates search for current end of data records.
FIND EVENT MARKER n	Searches for event market "n".
FIND TIMECODE xx.xx.xx.xx	Searches for time xx.xx.xx.xx.

Table 2 - Typical Status Set

(Penny & Giles PEGASUS)

STATUS REQUEST	RESPONSE
Immediate Status Requests	
?DATA INTERFACE	Lists all date interfaces currently installed.
?DATE	Date currently in calendar.
?ERROR	Current block error count.
?HOURS RUN	Current system hours run.
?MACHINE STATUS	Current status of recorder:-
	DATE year.month.day
	TIME hours.minutes.seconds
	MACHINE ID number
	TAPE LOAD STATUS status
	LAST COMMAND STATUS command
	DATA INTERFACE STATUS interface
?MEDIA STATUS	Current status of media:-
	TIMECODE xx.xx.xx.xx.xx.xx
	MEDIA ID number
	DATA INTERFACE type (recording).
	x% REMAINING TAPE TO EOT (End of Tape).
?SERIAL NUMBER	Recorder's unique ID.
?TAPE LOAD	Load status.
?TIME	Current system clock.
?TIMECODE	Most recent timecode read from tape.
Deferred Status Requests	
?EVENT MARKERS	Scans tape for and lists all recorded event markets and associated user data.
?MEDIA ID	Current media ID.

Figure 5. Data Flow and Hand-shaking.

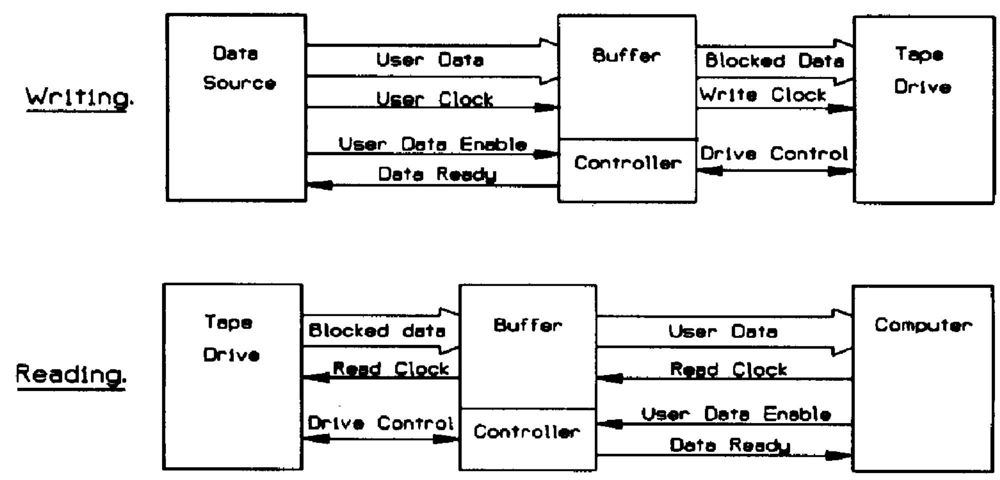

On replay, the situation is slightly different since it should be possible for the computer to control the transfer of data in accordance with its own needs and activities. Here, a hand-shaking protocol is essential since the mere fact that the computer may have requested data does not in every case mean that data will be immediately available. Consider the situation where a new cassette has been loaded into the transport and placed at the beginning-of-tape (BOT) but no other tape movement has yet taken place. The computer may request data and offer an output clock, but the recorder's buffer as yet contains no data. Instead, the recorder will acknowledge the request for data and immediately start to move tape in order to fill the buffer. At a certain point, there will be sufficient data within the buffer for an output transfer to commence. As long as the computer continues to demand data, the recorder will maintain an appropriate level of data in its buffer, starting and stopping the tape as necessary. At some point in the transfer process, the computer may decide that it has sufficient data and cease to request further data. Recognizing this, the recorder will discontinue the reading process although some valid data may remain in the buffer ready for transfer later.

A convenient method of achieving this is to use a common, bidirectional data input/output interface including hand-shaking lines which control the flow of data to and from the recorder. For example, a DATA READY signal may be asserted by the recorder to indicate that it is ready to receive data and a USER DATA ENABLE may asserted by the user to indicate that applied data is valid. When reproducing, a DATA READY signal asserted by the recorder means that valid data is available, while USER DATA ENABLE is asserted by the user to indicate that he is ready to accept outgoing data.

A PRACTICAL IMPLEMENTATION

An important new entrant in the race towards computer-friendly data recorders is the PEGASUS High Rate Digital Cassette Recorder from Penn & Giles Data Systems.

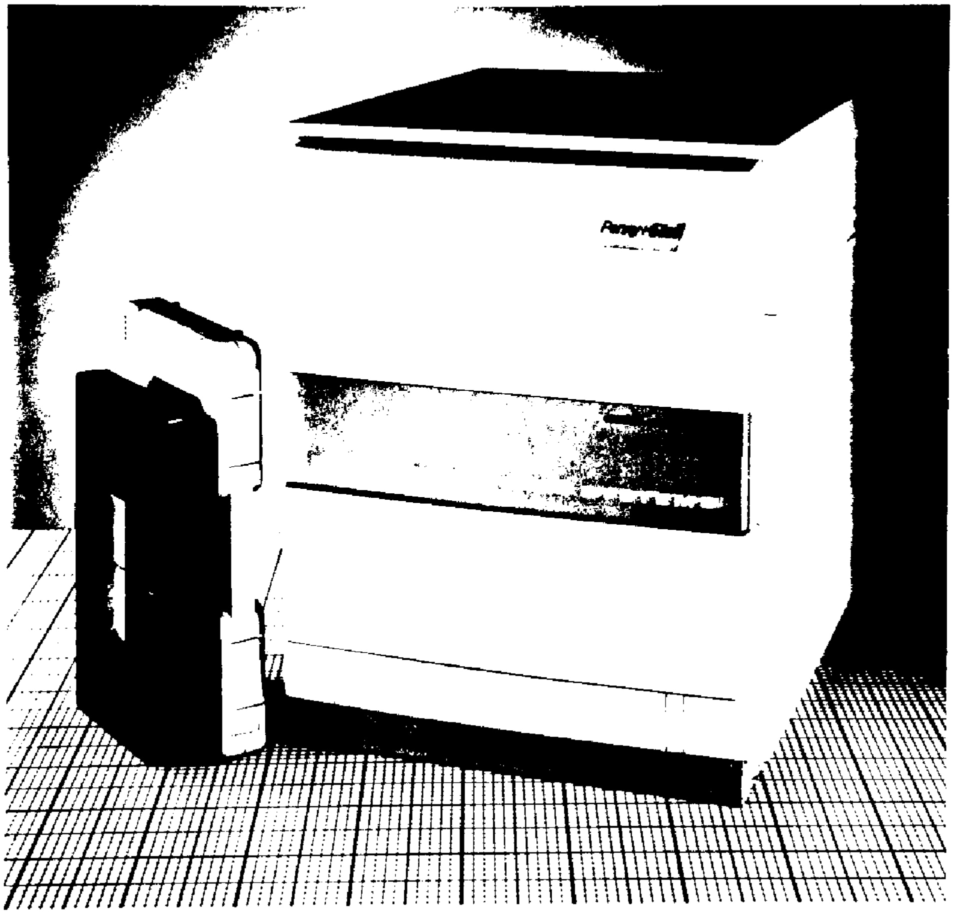

Figure 6 PEGASUS High Rate Digital Cassette Recorder

The following is a brief technical overview which may be helpful in summarizing the points which have been previously discussed.

Figure 7 shows the data block diagram. Data may be input or output in either parallel or high rate serial form via the data interface. The maximum continuous data rate for the recorder is 100 Mbit/sec (12.5 MByte/sec) with a burst rate up to 120 Mbit/sec (15 MByte/sec). Following any necessary format conversion (such as serial-to-parallel), a 32-bit wide data word is passed through the 256 Mbit data buffer under the direction of the recorder's system controller and into the encoder, parity generator and CRC generator. Here, auxiliary data such as event marking, time code and internal data block addressing are added before the combined data and control information is written to tape incrementally at a fixed tape speed.

On replay, data is read from tape incrementally then equalized, bit-synchronized, decoded, deskewed and error-corrected before being clocked into the data buffer. Any necessary format conversion (32-bit to 8-bit or parallel-to-serial) is accomplished in the user data interface prior to the data being clocked-out on demand. Auxiliary data is routed to the recorder's system controller via a separate path as shown.

The control of data flow into and out of the recorder operates at three levels: -

1) The main recorder functions and status (WRITE, READ, STOP, etc.) - via the control port.
2) File management functions (event marking plus time and event and searching) - also via the control port.
3) Data flow hand-shaking - via dedicated lines within the user data interface.

Figure 7. PEGASUS Data Block Diagram

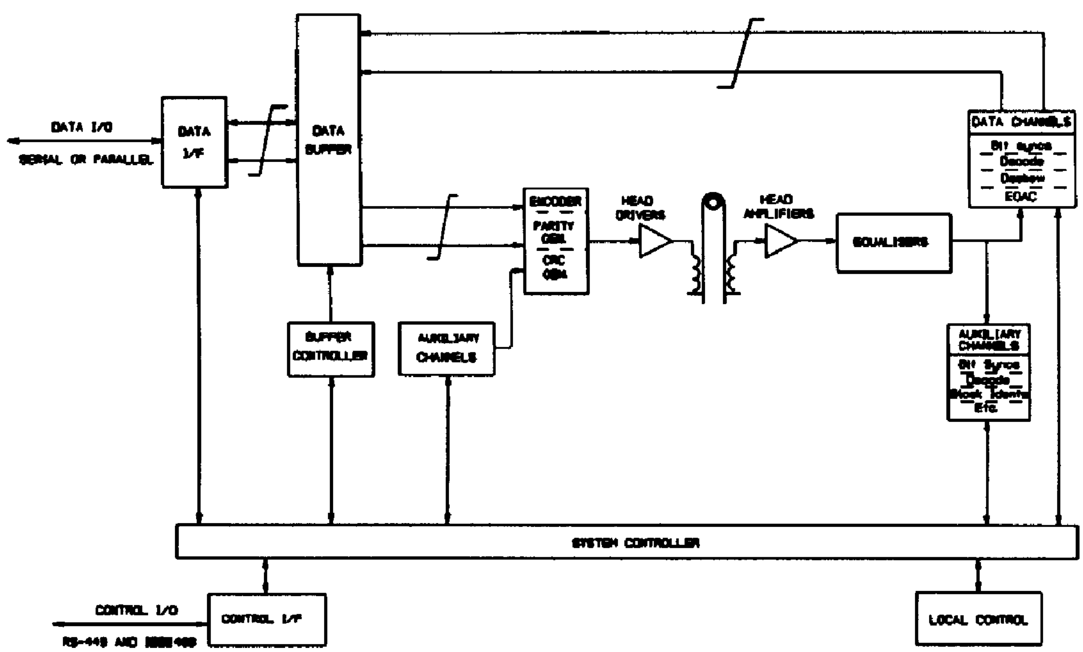

CONCLUSION

Recent market entrants have shown that it is now indeed possible to combine the merits of traditional high rate, high capacity data capture systems with the requirements for convenient and straightforward computer interfacing. The operational capabilities of this new class of cassette-based equipment have far-reaching implications for the way data will be captured and processed in the future. To take full advantage of the opportunities now presented, a critical assessment of traditional methods may be necessary. Indeed, we may soon see the day when the cultural divide between "instrumentation recording" and "computer storage" has been bridged forever.

Table 3 - Summary

Traditional Instrumentation	Computer Friendly
Continuous data transfer	Fast start/stop data transfer
Data 1/O buffering optional?	Integral data 1/O buffer
Complex computer interfacing	Simple computer interfacing
Separate data & control interfaces	Optional combined data & control interfaces
Data loss during ramp up/down handshake	Instant-on data write/read with data flow control
Slow tape load/unload time, must rewind before unloading	Fast tape load/unload time, rewinding unnecessary for unloading
Control of tape drive: RECORD, FORWARD, REVERSE, FAST, STOP	Control of data transfer: WRITE, READ, FIND, STOP DATA
Optional time code generate/read/search	Data management including integral calendar/clock & event generate/read/search
Data transfer rate proportional to tape speed	Data transfer rate controlled by user read/write clock and independent of tape speed