DIGITAL WIDEBAND RECORDING SYSTEMS

Alfred N. Brower
President, DSP*Con*, Inc.

OVERVIEW

Historically, those interested in recording one or more channels with analog content of greater than or equal to 2 MHz, must use an analog recorder. In the last few years, advancements in analog-to-digital converter technology, performance enhancement in Digital Signal Processors (DSPs), and digital recording devices have made cost-effective, wideband recording applications possible through the use of all-digital techniques.

This paper has three objectives:

1. It attempts to explain the benefits of a wideband digital recorder over the traditional analog variety.
2. It discusses the key elements of a wideband digital recorder.
3. It presents a realizable 10-channel, 30 Mbit PCM digital recorder solution.
4. It presents a realizable 14-channel, 2 MHz (bandwidth) digital recorder solution.

ANALOG VERSUS DIGITAL RECORDING

There are several advantages to digital recorders over traditional analog recorders. This section describes some of these advantages.

RECORDING EFFICIENCY

Typically, an analog recorder consists of a tape unit that is capable of placing each channel on a slice of the tape called a “track”. The greater the number of channels, the greater the number of tracks. Since all of the channels are written on the same tape, the tape moves at a speed capable of recording the highest bandwidth signal. If other channels contain lower bandwidth data or no data at all, the, in effect, this tape is wasted.

Using a digital approach, real-time processors are capable of packing digitized data in such a way that no part of the media is wasted. This allows the operator of the recorder to specify the maximum analog content on each channel; the recorder automatically performs in an optimal manner. Similarly, digital devices are also capable of taking advantage of situations in which some number of input channels are not being used (i.e., no signal is being applied). In these cases, a digital recorder ignores the inactive channels, thus allowing more data to be recorded on the active channels.

LONGEVITY OF DATA

With an analog recorder, no two recordings of identical input data will be recorded in exactly the same way. Likewise, when data is played back from a given analog recorder, its content will vary ever so slightly. This effect is even worse when data from one recorder is played back from another. All of this is due to the fact that the reproducibility of the true signal is based on how close to the intended speed the tape is moving, how much side-to-side movement the heads are producing over their intended tracks, and how clean the heads on the recorder are.

With a digital approach, all of the variance is located in the sampling clock of the record and reproduce systems. These are the clocks that feed A/D converters on the record system and the D/A clocks on the reproduce system. Variances of sampling clocks between systems result in variance of the frequency content of the data being captured and replayed. Today, inexpensive oscillators are available that are stable to more than 1-part-per-MHz in a wide temperature range. This results in a frequency stability better than any analog recorder, and far better then is practically required.

RELIABILITY

A major problem with analog recorders is that it takes a highly-trained user to operate one in a mission-critical environment. If one or more channels fail, it is difficult to tell until it is too late. Analog recorders suffer failures that are analog in nature.

Digital recorders are capable of providing very robust real-time error protection (e.g., Reed Solomon Coding). Also, they are capable of performing real-time write-read-verify cycles to all of the data they record. This way, the recorder can alert the operator as soon as a problem occurs. In addition, most digital recorders are capable of End-to-End (E to E) testing. In E to E testing, a signal of known frequency content (e.g., a sine wave) is self-injected into the A/D front end of the recorder. The recorder digitizes this data, writes it to tape and then reads the data back and ascertains that the spectral content of the recorded data is adequate to allow the user to continue. Lastly, digital recorders always have failures that are digital in nature. In other words, when they work, they are at 100%. When they fail, the operator is immediately informed and it is impossible to continue. Often, this is very desirable--particularly in mission critical environments; because we are operating on the data as it is coming in, we can perform a number of real-time checks on it simultaneously.

MAINTENANCE

Anyone who has used an analog tape recorder is familiar with the expense of having the tape heads replaced. Typically, this is not an option because failure to take timely action results in degraded data or data that is lost completely. In addition, analog recorders, in general, require the operator to make periodic adjustments to insure proper tape alignment, proper tape tensioning and set up of record/reproduce electronics.

With a Digital recorder, there is virtually no maintenance. For example, disk-based systems have a mean time between failures of nearly 10 years. This number is even better for digital tape systems, although these systems do require periodic head cleaning, this is easily done with a cleaning tape.

REAL-TIME PROCESSING

Most analog recorders provide no real-time analysis to the operator concerning the data that is being recorded. Even the most sophisticated recorders provide little more than limit checks, time series and, perhaps, spectral viewing of the data.

Once the data is digitized (which is the front end of every digital recorder), today's enormously powerful Digital Signal Processors can provide a wide variety of real-time feedback to the operator. This includes a wide variety of PSD and Spectral views (with selectable windowing), filtering, demodulation, channel-to-channel correlations, and signature analysis and recognition (including database searching).

DATA FORMAT VERSATILITY

Typically, the data on an analog recorder is physically stored on the tape in one format or another. This format is determined at the time of purchase. In general, data from one vendor's recorder may be replayed on another's as long as the format is identical. As the number of standards grew, all to often consumers found themselves making a difficult decision to buy new recorders or live with a outdated data format for which fewer and fewer recorder vendors provided support.

The heart of all DSP*Con* digital recorders is a very fast processor, which is not only capable, of optimally packing the data, but also of formatting the data in a variety of standards. As standards change, DSP*Con* digital recorders can be updated with a simple firmware modification, which is certainly less expensive then replacing the entire recorder.

HOW DIGITAL WIDE BAND RECORDERS WORK

Most DSP*Con* digital recording systems consist of the same three basic blocks. They are:

- Filtering and Analog to Digital Conversion
- A Digital Multiplexer and Logical Data Formatter
- A Disk Array or Digital Tape Drive

A diagram of these important blocks is shown in Figure 1.

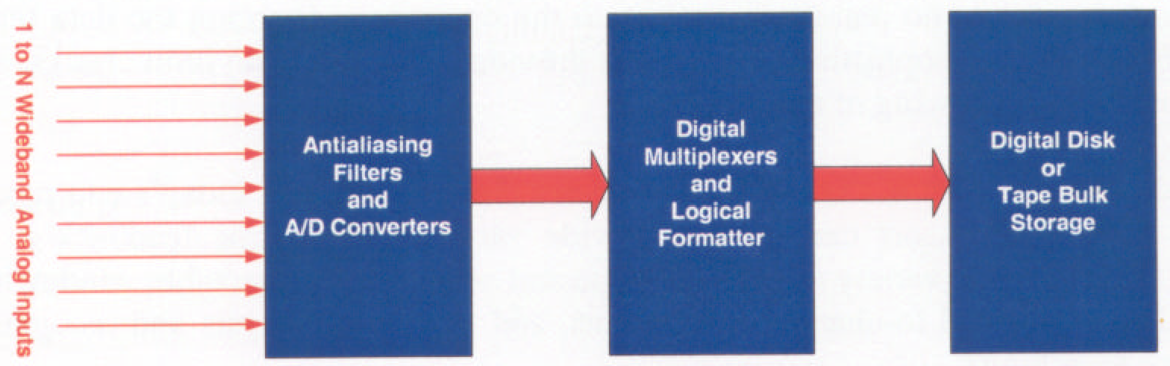

Figure 1-Basic System Blocks

Please note that the number of input channels can be traded off against the input bandwidth (sampling rate) as long as the aggregate system data rate (number of channels multiplied by the digitized byte rate of each channel) does not exceed the bandwidth of any bus in the system or the input bandwidth of the bulk storage system. For example, DSP*Con* has delivered 4-channel systems with 30 MfIz of input bandwidth for each channel. DSP*Con* has also delivered 96-channel systems, with a very similar architecture, in which the input bandwidth was 45 kHz.

This section discusses each of these blocks and describes the state of the art technology of each.

FILTERING AND A/D CONVERSION

Arguably, this is the most important part of the recorder. It is the only part of the recorder that handles the analog signals that will be captured. The first part of the block is responsible for removing spectral content on the signal, which is beyond our frequency of interest. This step is necessary to prevent an effect known as “aliasing”, in which these higher (unwanted) frequencies may corrupt the signal of interest as the data is being digitized. While a discussion of digitizing techniques is beyond the scope of this paper, it is a very important part of the front end of a digital recorder. A system that lacks proper anti-aliasing filters results in the recording of corrupted signals. It is important to note that this filtering must be done with analog components. There is no technique for removing the effects of aliasing once the data has been digitized.

Once the analog signals have been sent through the proper anti-aliasing filters, the data is ready to be converted into digital form. The concept here is actually quite simple. A/D converter chips are fed with a clock. Each clock edge causes the converter to quantize the analog level that is present at the converter at the time of the clock edge. The faster the provided clock, the more often the A/D quantizes. In order to capture all of the desired spectral content, one must sample (clock the A/D) at a theoretical minimum rate of twice the maximum spectral frequency of interest-known as the “Nyquist rate”. For example, if we are interested in recording all analog content up to 2 MHz, we must sample at the rate of at least 4 MHz. As actual silicon does not behave at exactly the Nyquist rate, most engineers sample at a minimum sample rate of 2.5 times the maximum frequency of interest. This accounts for imperfect filter roll-off and other realworld effects. While many engineers believe that sampling at rates much higher than the Nyquist rate is required in order to capture all of the spectrum of interest, a DSP*Con* White Paper entitled “*Interpolation of Sparse Time History Data*” explains why this is not the case. It can be found in its entirety on the DSP*Con* web site (www.dspcon.com). To access the paper, choose “Download” from the menu choices, then choose “Tech Papers” and scroll to the document.

By quantizing (sampling) often enough, we can assure that we capture all of the spectral content of interest. Since we are using a finite number of bits in each sample to characterize an infinite number of points on each analog signal, how accurately can we capture a complex analog wave? The answer is driven by the number of bits in each sample. This is measured in a ratio of the actual signal to noise (or error) in the sample. This ratio is called the SNR (Signal-to-Noise Ratio) of the system. Years ago, if you wanted to sample at 2 MHz, you would be “stuck” with 8-bit samples that resulted in an SNR of about 42 dB. Today, it is very common to sample 2 MHz signals with 16-bit A/D converters, which yields a SNR of almost 80 dB. Compare this to the less than 30 dB achieved from any reel-to-reel analog recorder. This performance is far better than any analog recorder that has ever been made and usually far better than is required (given the noise associated with sensors, cabling, and other parts of the system).

One final note about the front end of the recorder. Analog recorders maintain a poor time correlation between each channel that is being recorded. Static and dynamic inter-channel time displacement errors create many microseconds of time discrepancy between tracks. Typically, DSP*Con* recording systems maintain a channel-to-channel time correlation better than .1 degrees at 1 kHz.

DIGITAL MULTIPLEXER AND FORMATTER

Assuming that we are interested in 14 channels and 2 MHz analog bandwidth, we must sample at 5 MHz. In addition, assume that each sample contains 16 bits (2 bytes). This results in 140 Mbytes/second. As little as two years ago this would have overwhelmed even the most powerful processors. Today, several processor vendors, for example, Intel, Texas Instruments, and Motorola, are shipping processors capable of billions of operations per second. These powerful processors are able to manipulate this amount of data all in real time. At a minimum, this block must input all of the data from the A/D front end, and supply this data as a single time-packed stream to a disk or tape device. Better systems can actually analyze this data and provide useful statistics to the operator, while also streaming the data to the bulk storage. Arguably, the most important task that this block should perform is logical formatting of the data that will be written to the storage. This logical format allows end users to read the data and process it with algorithms that they may have purchased or developed. Typically, such formats include a specification of the number of channels being recorded, the rate at which they are being recorded, and time tagging of the data plus significant events.

One of the most important aspects of the digital multiplexer is its ability to process the data in addition to formatting the data. This offers a powerful solution to a large problem. The problem is how to manage the enormous quantity of data that is collected and needs to be archived by wide band recorders. Most often, these types of recorders are used to record long periods of data in order to be sure that specific events of interest are recorded. The powerful real-time data processors within a digital recorder offer two powerful ways to address the data reduction problem.

- Users may annotate blocks of interest with data marks that include long, text-based comments that can be searched for off line. As these marks can be searched for rapidly on either disk- or tape-based systems, data reduction can be performed quickly and inexpensively.

- The real-time data formatted can be programmed to search automatically for characteristics of signals of interest in the time and/or frequency domains. This applies to both the start and end of these events. Using this technique allows the recorder to record for much longer periods of time because only events of interests are recorded. Moreover, when a recording mission is concluded, there is only a small fraction of the data that needs to be archived.

Potential users of these all-digital, wideband recording systems must understand that planning is necessary to allow for recovery of specific data in a time-efficient manner. As the actual recording bandwidth per inch of media is increased by analog to digital conversion, users will not have the luxury of playing back data at increased speeds. There must be a tight firmware link between the real-time indexing of important data as it is recorded.

BULK STORAGE

There are two ways that bulk storage systems can be characterized. The first is the interface standard that is used to get data in and out of the unit. The second is the type of physical media within the unit that is used to record the data.

MEDIA TYPES

There are basically three different types of digital storage system media. Each has its benefits and its drawbacks. All, however, are digital and therefore offer all of the benefits detailed in the first part of this paper. Cost factors vary widely with each type of system, depending on environmental factors. As the military or intelligence markets no longer drive recording technology, systems tailored to those markets can be inordinately expensive. DSP*Con* strives to use COTS products wherever practical, but can use any digital medium which best fits the application.

- **Digital Tapes**—These devices offer the obvious benefit of having removable media. They also can be flown at relatively high altitudes and still work well. Lastly, most modem tape devices hold large quantities of data up to 660 Gbytes per tape. One large disadvantage of tape devices is that they tend to be slow. Even the faster tape units are capable of handling only 24 Mbytes/second.

- **Solid-state Devices**—Basically, a solid-state bulk storage device is a large amount of RAM configured to look like a disk or tape. These devices tend to be very mechanically robust and extremely fast. They do, however, also tend to be very expensive and are generally not removable.

- **Disk-based Systems**—Disk-based storage systems are made from piecing together two or more individual disk drives to produce a higher-capacity device, a faster device, or both. Disk-based storage systems tend to be fast and inexpensive. In addition, depending on how the disks are physically mounted with the storage device, they can also be very mechanically robust and even removable (like tapes). There are basically two types of disk-based storage systems.

  JBOD—A JBOD (Just a Bunch Of Disks) is two or more individually-addressable disks within a box. This is roughly equivalent to hooking up multiple disks to one interface. These devices are extremely inexpensive and offer large capacities as a result of adding together the individual capacities of all of the disks contained within the JBOD. Since many of today’s disks are capable of speeds close to 40 Mbytes/second and 72 Gbytes/disk, a 40 Mbytes/second, 360 Gbyte (5-disk) JBOD can be purchased very inexpensively.

  RAID—The problem with a JBOD is that while it allows you to add the capacities of each disk together, the slowest disk device within the JBOD determines the speed of the entire device. This is due to the fact that each disk is individually addressed for each write of data. A RAID is a variation on this theme. A RAID is also made up of a bunch of disks. However, unlike a JBOD, which writes its data to each disk in series, a RAID contains a piece of hardware that spreads each data byte across multiple disks. This allows us to add not only the capacities of each disk device, but also the individual speeds. If each disk is capable of 40 Mbytes/second, when five such devices are put inside of a RAID, the resulting speed is 200 Mbytes/second (5 times 40 Mbytes/second). Typically, RAIDs also contain an additional disk used to store the parity of each transfer. The benefit of this is that each data transfer to the RAID is read back and the parity is checked against the expected parity. If an individual device fails, that device will be immediately replaced with the Parity disk. As a result, RAIDs tend to be one of the most reliable types of bulk storage devices. Today, it is common to see RAID devices that have a capacity of 320 Gbytes, and are capable of write rates in excess of 200 Mbytes/second.

- ***Emerging Technologies***—As the past several years have shown, there is a good deal of innovation still taking place in data storage technology. These advancements include new media types, cost reductions (for both the recorder and the media), increases in recording rates, and significant increases in recording densities. In fact, over the past several years, recording densities have grown at the rate of approximately 30% roughly every 6 months. An important consideration to end users is how to stay current with technology trends without assuming the overwhelming financial burden associated with replacing all of the fielded recorders and dubbing all of the archived data from the outgoing media to the new media.

  Regardless of how one chooses to handle this problem, there is always the a cost associated with staying current with technology—digital recording systems tend to make this outlay less traumatic. This is because digital recording systems see bulk storage simply as a peripheral that hangs off of an industry standard interface. This is true regardless of the type of storage. Therefore, the storage device itself can get bigger, faster or even change its nature and not impact the recording system. This very important fact is also true for replaying archived data. As long as the data is digital, and the storage device has a common digital interface supported by recorder/replayer, a variety of media types can be archived effectively. DSP*Con* has had a great deal of success in swapping digital tape recorders with RAID, JBODS and vice versa. Moreover, DSP*Con* fully intends to continue to support new devices as long as these devices are digital and obey open interface standards.

The next section of this paper discusses some popular 1/0 to storage interface standards.

INTERFACE TYPES

Today, there are two different (albeit similar) methods of moving digital data to and from digital bulk storage devices. The fact that these are standards is very important because, in theory, this makes replacing one unit from one vendor with another unit from perhaps a different vendor easy. Moreover, in theory, upgrading a unit to the next generation should be relatively inexpensive. Each of these methods is described below.

- ***SCSI***—A SCSI is basically a parallel bus with address lines that can attach to up to fifteen devices. A processor (called a host) may initiate a transfer of a block of data to or from any device on the bus. A SCSI Ultra Wide bus is 32 bits and has a maximum sustained transfer rate of 40 Mbytes/second. SCSI buses come in two electrical types—*Single-ended* and *Differential*. The maximum length of a differential SCSI bus is approximately 20 meters--far less for single-ended. Please note that for each type of device that can be on a SCSI bus (fixed disk, removable media, etc.), there exists a very well defined and rigid protocol for communication with these devices.

- ***Fibre Channel***—Unlike SCSI buses, Fibre Channel is a serial interface that can be run up to 10 kilometers. It will support up to 1 Gbit/second. At the time of this writing, several vendors are already starting to ship 2 Gbit/second devices. There is no practical limit to how many devices can be present on a Fibre Channel loop. Most importantly, Fibre Channel supports two well-known protocols; SCSI and TCP/IP.

  This turns out to be an incredibly important fact. It means that Fibre Channel can be used for networking computers together and/or for interfacing computers to bulk storage. There are two major reasons why this is significant. The first is that since Fibre Channel is a networking standard, this technology is well supported by hundreds of vendors, and that peripherals, including switches and 1/0 interfaces, are inexpensive. In addition, there is a great deal of commercial industry focus for enhancing this technology. The second reason is that people are now building standard networks that consist of processors and bulk storage devices. As long as the bulk storage devices have a Fibre Channel interface, they appear to the network as simply another client device. This is important because it means that the wideband recorder may readily be connected to many back office networks for data extraction.

APPLICATION EXAMPLE ONE-A 10-CHANNEL, 30 MBIT PCM DIGITAL RECORDER

This system acquires up to ten asynchronous channels of 30 MBits/second PCM data via the DSP*Con* model 9100VME digital 1/0 card, which then passes the data to a DSP*Con* Model 9245 DSP baseboard where the data is processed and streamed to a Fibre Channel PMC module. Finally, this PMC module routes the data to a fibre channel switch, which communicates with a workstation and records the data to a Sony GY-2120, GY-8240 DTF or DTL-7000 recorder. Data can be extracted from the recorder and routed back out to the source with up to .001 Hz of the recorded input rate.

This system provides the following features:

- Accepts 2, 4, 6, 8 or 10 inputs.
- Each input may be differential RS-422 or TTL.
- Each input is capable of up to 30 Mbits/second.
- IRIG Time Code (IRIG A, B or E).
- A single voice channel.
- All I/O is provided via BNC connectors at the rear of the unit.
- All data is archived to digital tapes (DTF). Since data is stored digitally, the unit only records data on channels where data is present.
- The unit automatically reacts to data suddenly, appealing or disappearing on channels.
- Storage is scalable; 200 Gbits, 336 Gbits or 1.6 Tbits (combined across all channels).
- Operator controls the unit via a user-friendly Graphical User Interface (GUI), on a standard PC, which provides the following controls:
  - Start record
  - Stop record
  - Inject Tape Mark
  - Inject Comment
  - Search to Tape Mark
  - Search to time Code
  - Eject Tape
  - Play
- The GUI also provides the following status:
  - Tape Remaining
  - Tape Consumed (as a percentage)
  - Time Remaining
  - Data Rates on Each Channel
- Optional bit syncs.
- Optional replay only capability.
- Optional upgrade capability to comply with IRIG 107

A block diagram of a solution to this application is shown in Figure 2.

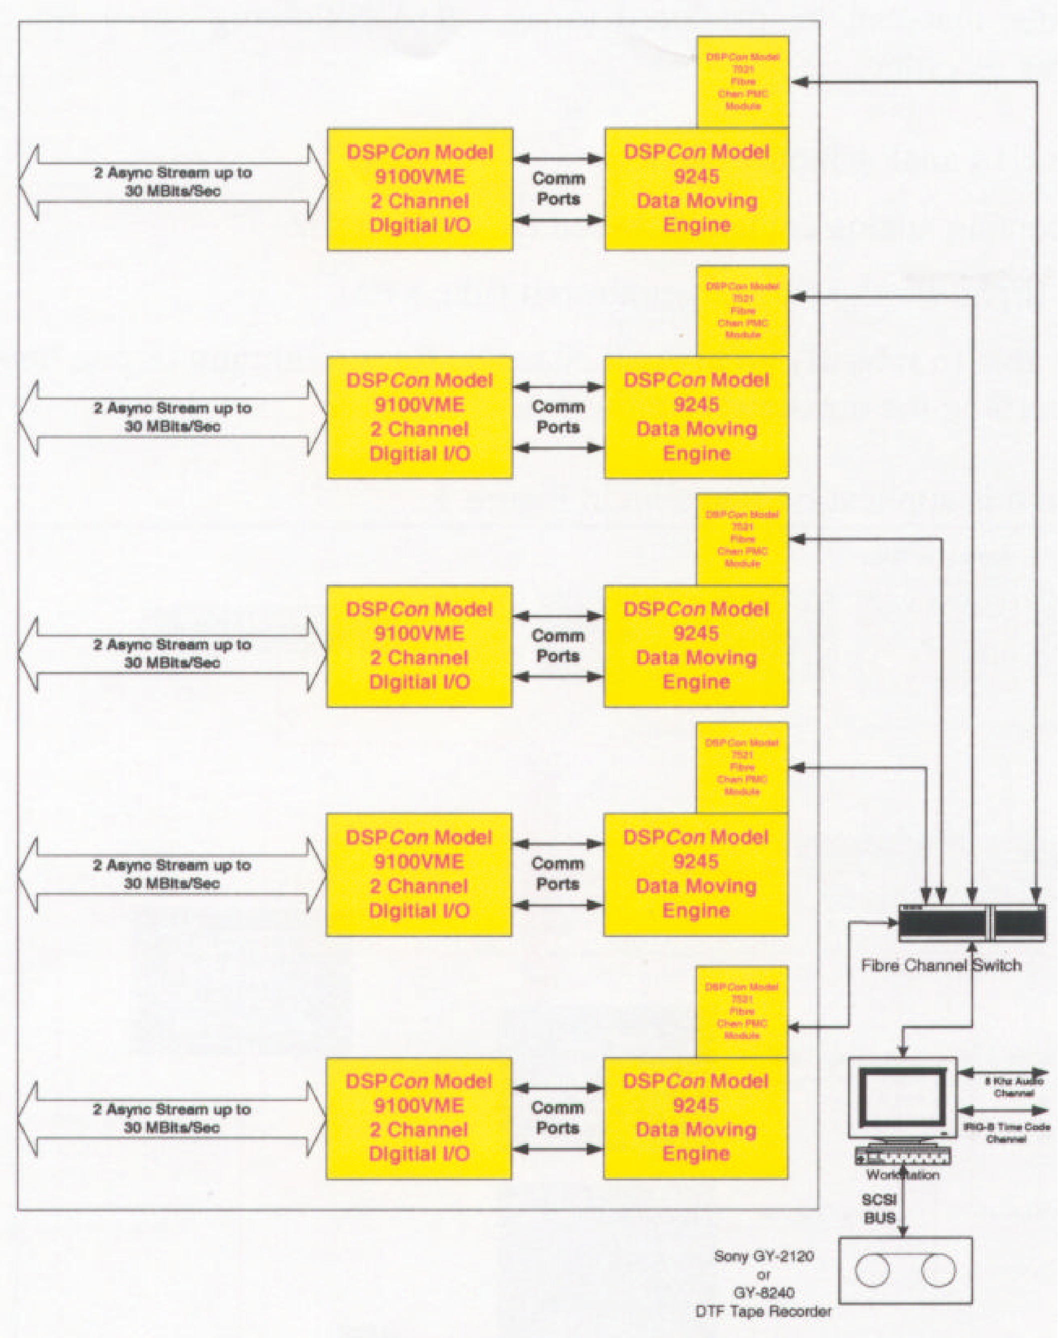

Figure 2 – System Block Diagram

APPLICATION EXAMPLE TWO ¾ A 14-CHANNEL, 2 MHZ WIDEBAND RECORDER

Using the ideas discussed in this paper, we will present a block diagram of a cost-effective, multi-channel wideband recorder that can be produced today. The following are a list of hypothetical requirements for this recorder.

1. The system must have 14 analog input channels.
2. Each channel must contain analog content between DC and 2 MHz.
3. The system must be flight-worthy on a larger aircraft (like a P3).
4. The system must be able to robustly archive all channels for a minimum of one hour, with all channels recording the maximum bandwidth.

A block diagram of a solution to this application is shown in Figure 3.

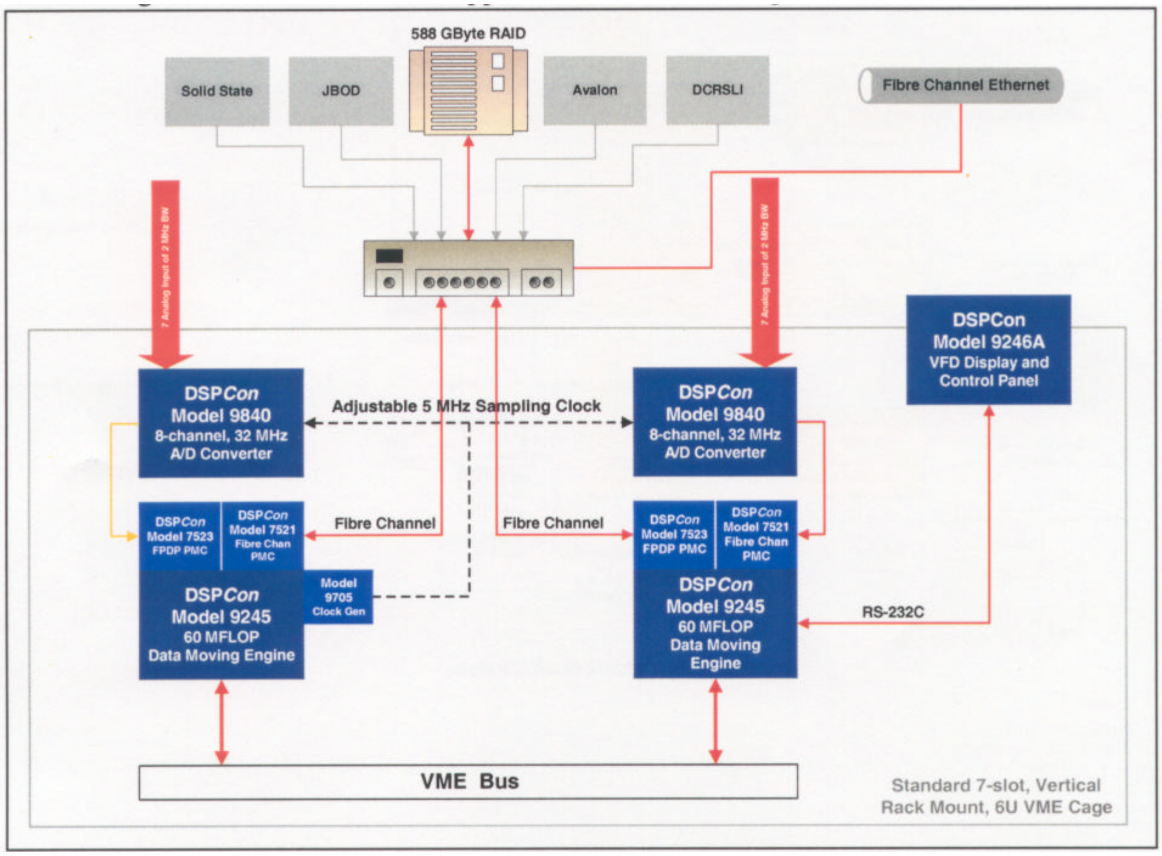

The system consists of three physically-separate blocks, each of which are rack mountable in a standard 19-inch rack. Each of the three sections of a standard wideband recorder (as presented in this paper) are discussed below.

FILTERING AND A/D CONVERSION

Each of the signals of interest are routed from the rear of the box to an input connector of one of two DSP*Con* Model 9840 A/D cards. The 9840 is an 8-channel, 16-bit, 35 MHz (maximum per channel) A/D card with input filtering. This product consumes two, 6U VME slots (per 9840). The first slot contains all proper analog filtering necessary to remove all aliasing beyond Nyquist. The second half of each 9840 contains all of the A/D converters (8 per card), digital packing logic, buffering and sample clock control. The 9840s output their data (only 7 channels in this case) to an industry-standard bus called an FPDP (Front Panel Data Port). Lastly, please note that the 9840 is capable of accepting and external input clock. This is important because there is more than one A/D card in the system, and we want to assure ourselves that all channels are being synchronously sampled.

A data sheet on the DSP*Con* Model 9840 has been provided in an appendix to this paper.

DIGITAL MULTIPLEXER AND FORMATTER

Each DSP*Con* Model 9840 presents its data via an FPDP interface to an FPDP interface model on a DSP*Con* Model 9245. The 9245 contains a high-speed DSP card that is capable of three concurrent activities.

1. Using its advanced mathematical core to process the data in real time.

2. Archiving the raw and/or processed data via a Fibre Channel interface module on each Model 9245.

3. Providing the data to the Fibre Channel interface on which a standard host may be attached for real-time display of the data. Please note that this is shown as an Ethernet connection in the block diagram.

A data sheet on the DSP*Con* Model 9245 has been provided in an appendix to this paper.

BULK STORAGE

In order to meet the 140 Mbytes-per-second I/O requirement and the one-hour record duration, a 588-Gbyte RAID has been selected. The size of the RAID is determine as follows

14 Channels * 5 Msamples/SecondlChannel * 2 Bytes/Sample * 3600 Seconds = 504 Gbytes

Please note that all devices are connected together by a low-cost Fibre Channel Hub. This allows additional protection between the devices if one device were to fail.

READING THE DATA FROM THE RECORDER

There are two ways that data may be “removed” from this system. The more cost-effective of the two is to hook one or more standard host computers (PC, Sun, HP, etc.) up to the Fibre Channel Hub and simply read all of the data from it. This may be done at rates up to 200 Mbytes/second.
As an alternative, RAIDs with removable disk packs can be provided. These disk packs come out as a group and new disks can be inserted in their place (much like a standard tape device).
Please note that there is no reason why both of these techniques cannot be applied together and, depending on the operational scenario, one method or the other can be chosen.

CONTROLLING THE RECORDER

The entire recorder may be controlled in one of two ways.

1. The 9246A provides a simple push-button display that allows the operator full control of the recorder mode (Record, Stop, Forward, Rewind, Number of Channels to Record, Sampling Rates, etc.). A display indicates the mode of the recorder, BIT Status, the amount of data currently recorded, the amount of space remaining on the recorder, the number of channels being recorded, and the sampling rate of the A/D.

2. All of the control and status functions provided by the 9246A (above) might also be provided via a PC host computer connected to the Fibre Channel. This has the benefit of also allowing a large number of real-time, graphical displays and analysis of the data.