



Fermilab

TM-1363
2110.000

A REAL TIME DATA COMPACTOR (SPARSIFIER)
AND 8 MEGABYTE HIGH SPEED FIFO FOR HEP*

A. E. Baumbaugh, K. L. Knickerbocker, and C. R. Wegner
Fermi National Accelerator Laboratory, Batavia, Illinois

and

B. W. Baumbaugh and R. Ruchti
University of Notre Dame, Notre Dame, Indiana

October 1985

*Submitted to the 1985 IEEE Nuclear Science Symposium (NPS),
October 23-25, 1985, San Francisco, California.

**A REAL TIME DATA COMPACTOR (SPARSIFIER) AND
8 MEGABYTE HIGH SPEED FIFO FOR HEP.**

A. E. Baumbaugh, K. L. Knickerbocker and C. R. Wegner
Fermi National Accelerator Laboratory,*
Batavia, Illinois
B. W. Baumbaugh and R. Ruchti
University of Notre Dame, Notre Dame, Indiana

Abstract

A Video-Data-Acquisition-System (VDAS) has been developed to record image data from a scintillating glass fiber-optic target developed for High Energy Physics.[1],[2] The major components of the VDAS are a flash ADC, a "real time" high speed data compactor, and high speed 8 megabyte FIFO memory. The data rates through the system are in excess of 30 megabytes/second. The compactor is capable of reducing the amount of data needed to reconstruct typical images by as much as a factor of 20. The FIFO uses only standard NMOS DRAMS and TTL components to achieve its large size and high speed at relatively low power and cost.

Introduction

A block diagram of the overall system is shown in Figure 1. The active target and video data acquisition system are only a small part of a much larger experiment at Fermilab. The unique features of the VDAS discussed here are the data compactor and the 8 megabyte FIFO, both of which have data-throughput rates of at least 30 megabytes/second.

The scintillating glass fiber plate acts as an active target for the experiment and will be used for vertex detection and for identification of particle decays which may happen near the interaction vertex. The scintillating glass target produces an image of any charged particles that pass through or interact within it. The light produced by the target is at the single to few photon level per fiber and needs multiple stages of image intensification to be seen or recorded.

A video camera is used to convert the image into a form that can be digitized for input to a computer. (Early glass tests used a film camera.) The video control section keeps camera, ADC, and compactor electronics in rigid synchronization to assure that there is no pixel jitter. The data compactor operates on the flash ADC video data, compressing it to reduce the amount of storage needed to reconstruct the image. The compactor is constrained to cause no degradation of image information. Any loss of image resolution would make later reconstruction of the physics of the interaction much more difficult.

The FIFO memory is used to buffer the high speed bursts of data from the compactor allowing the computer to take the data at its own rate. The control CPU is used to move data from the FIFO onto mass storage in the form of an optical disk. The CPU also copies selected frames into a video display generator for viewing.

Active Target, Intensifier, and Camera: The active target consists of a coherent fiber-optic bundle made from scintillating glass. Scintillation light produced by charged particles is proximity focused onto the input Fiber-Optic-Plate (FOP) of a multi-stage image intensifier which gives enough gain to produce a visible image. This image is coupled to

the input FOP of a video camera. The camera may be a vidicon, a CCD or a SIT. The read out of the camera need not be at standard TV rates, i.e. the SIT can be read at 50 MHz pixel rate

Camera control: The video controller was made programmable to a maximum image size of 4096 by 4096 pixels and initially a maximum pixel rate of 30 Mhz., since the choice of camera and pixel rate have yet to be determined. (This rate will be increased to 100 Mhz in the near future.) The camera control can be triggered by the experimental trigger and will digitize one frame and stop. The initial camera used for testing the system was a vidicon, with a resolution of 640 by 256 pixels. A CCD camera, 384 by 244, was then used to record actual cosmic-ray events for testing different active targets.

ADC: The flash ADC used at this time is a Siemens 6020 running at a maximum rate of 30 MHz, it produces 6 bits of data corresponding to 64 levels of intensity or grey scale. The intensity values of the pixels will be used to reconstruct the tracking information to a higher resolution than is inherent in the resolution of the camera.

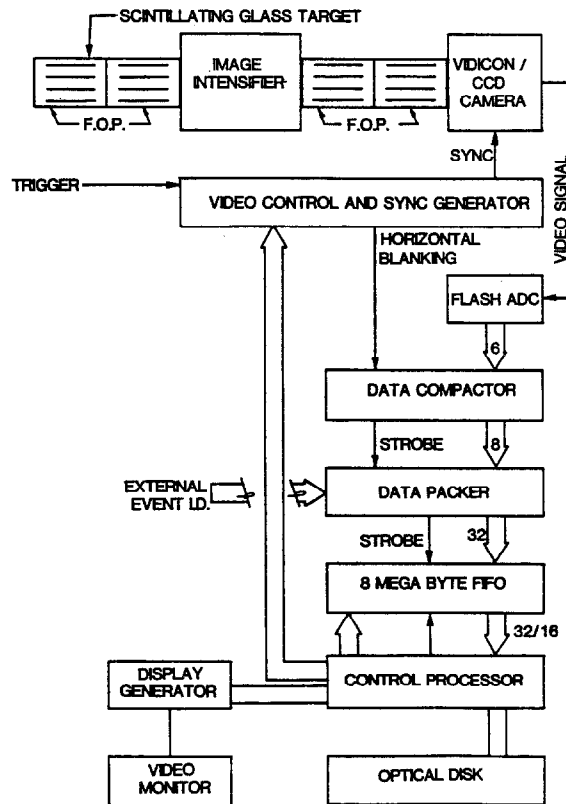


Fig.1 Overview of Video-Data-Aquisition system to be used in Fermilab Experiment E-687.

*Operated by Universities Research Association Inc., under a contract with the United States Department of Energy.

Data Compactor

The amount of data in an image can be very large. An image 512 by 512 takes 256kbytes to store in non-compacted form. Since typical images in this system have only a small percentage of their pixels above threshold, the image can be compacted to save storage space and time. (The maximum data rate through the entire system is limited by the rate at which data can be written to the optical disk. Compacting the data allows more images/second to be recorded.)

A prime requirement for the compaction scheme is that the image be totally reconstructable, since its used to reconstruct the interaction vertex. The data must be complete, unlike some image compression schemes where the result has only to be pleasing to the human eye. The compaction scheme is also required not to be "explosive", which is to say that, if a large number of pixels are above threshold, it should not take more memory to store the "compacted data" than the un-compacted data. The scheme used has both these properties in that the data is exactly reconstructable and if there are no zeros the space needed to store the frame is 1 byte/pixel.

Compaction Method

The method of compaction is to count and discard "zero"s (a "zero" is any value below a preset threshold) until the "zero" count equals 127 or a non-"zero" value is found. If a non-"zero" value is found the "zero" count is written, followed by the pixel value. If the "zero" count reaches 127 then that count is written and the "zero" counter reset. Only "zero" counts of one or more are ever written, so that consecutive non-zero pixel values are not separated by extraneous reset "zero" counts.

The compacted data word must be at least 1 bit larger than the input data words being compressed to allow for a flag bit in the data. The current system uses 2 bits for flag data which allows the inclusion of sync markers as fiducials in the data which simplify image reconstruction. There are three types of bytes in the compacted data stream, these are:

- 1) "Zero" count of the form 0XXXXXXX where XXXXXXX is the "zero" count 1 to 127.
- 2) Data which is in the form 10XXXXXX where XXXXXX is the 6 bit flash ADC data for a pixel.
- 3) Status or flags of the form 1100XXXX where XXX indicates the type of flag: 000 for a horizontal sync fiducial, 011 if 3 bytes of run number follow, 100 if external trigger ID follows, 101 if frame byte count follows, 110 if frame number follows, 111 if threshold value and status follow. These flag bytes and their data are used to re-link the image data, which is written to optical disk, to the remainder of the experimental data which is written to separate magnetic tapes during data taking.

It should be noted that while this compactor is used on video data it will work with any digital data stream and can be modified for virtually any width data word.

Compactor Hardware

While this scheme for data compaction is straight forward, the speeds at which the system must function pose a formidable task. The compactor must operate at the speed of the ADC which for these tests was set to a maximum of 30MHz. This implies that a new data byte is present to be compacted every 33 nanoseconds. A block diagram of the compactor is shown in Fig. 2. The output of the ADC is ECL and the compactor consists of ECL 10K series components in order to achieve the required speed. We believe that the use of ECL 10KH will allow for compactor speeds of 100 Megabytes/second.

The method used to achieve the required speed for this encoding scheme is that of data pipelining. All the latches shown in Fig. 2 are gated by a common strobe derived from the master pixel clock. Every 33 nsec. the data moves down 1 latch level. The pipelining of the data paths allows more than enough time at each level for the gate delays needed to make the decisions.

The first latch allows the maximum amount of time for the magnitude comparison. The second latch's output has both the data and the result of the magnitude test. The third latch adds to the data stream the "zero" count and status bits indicating if the "zero" count is 0 or if its 127. These status bits are used to decide what types of loads are performed later. The fourth and final latch gates through all the information needed to produce the load strobes for LDZ (load "zero" count), LDD (load data), LDHS (load horizontal sync marker), as well as "zero" counter control lines for INCZC (increment "zero" count) and CLRZC (clear "zero" count). All of these signals plus the data and "zero" counts appear at the output of the final latch at the same time. These signals are available for 33 nsec., during which the load strobes occur as needed. The load control logic operates a mux and strobe select to load the appropriate data. This entire system is designed to run synchronously with strobes continuously produced but only enabled under the proper conditions.

The 8 bit wide output of the compactor is packed into a 32 bit wide word and converted to standard TTL logic levels for transmission to the FIFO.

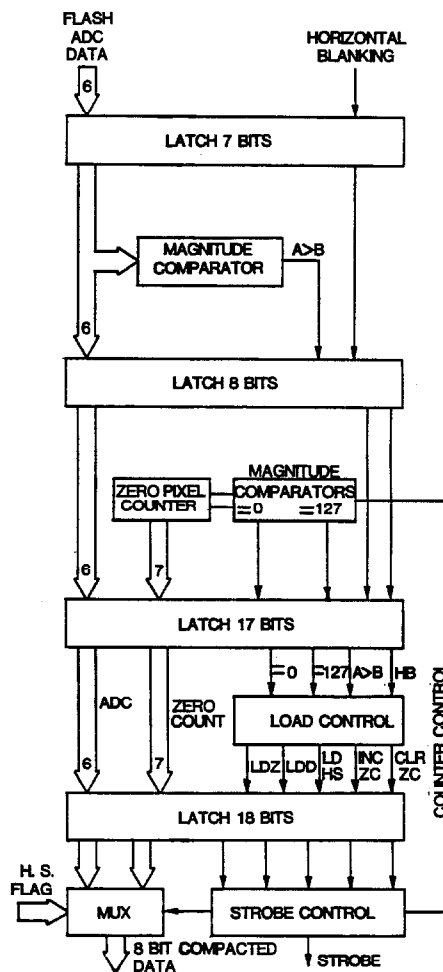


Fig.2 Block diagram of the data compactor, note pipeline of latches to allow speeds needed.

The maximum compaction factor for an 8 bit wide data stream which is all zeros is 127. In operation, discussed later, the compactor has achieved compaction factors as high as 30 for sample cosmic ray images.

Compactor Adaptability

While the video image data occurs at a uniform rate, the compactor could be gated externally with data that does not occur uniformly as long as the maximum clock rate was not exceeded. This would allow the digital compaction of a variety of data streams at very high speeds. In addition, it should be noted that this scheme of data compaction easily lends itself to hardware de-compaction. While we have not built a de-compactor, (the images are reconstructed off-line), one could be built which would allow data streams to be packed and un-packed in a way which is transparent to the user.

Eight Megabyte FIFO

Image data storage requirements for the VDAS are rather difficult in that large amounts of data are generated by the ADC/Compactor in a relatively short period of time. We have developed an eight megabyte FIFO that meets these demands while keeping in mind the ideals of low overall cost and low power consumption. This FIFO can in principle handle input data rates of up to 73 megabytes/sec, providing a smoothing effect between itself and the host computer, at this time however, no attempt has been made to exceed 30 megabytes/sec.

The FIFO uses ALS/AS TTL for input/output shift registers, storage registers, and all FIFO control circuitry. These components attain the required speed with low power consumption. The FIFO's eight megabyte main memory consists of 32 NMOS 256K x 9 DRAM SIP modules, chosen for low cost, low power, and ease of replacement. These choices allowed a substantial savings in power when compared to an equivalent amount of ECL ram. The ECL ram alone would dissipate approximately 15KW while this FIFO when running consumes only 32.5 watts RMS.

FIFO Organization

Externally, the FIFO appears 32 bits wide and 2 megawords deep. Internally however, the FIFO is organized as a 256 bit word 256K words deep. The system uses 128 ALS/AS TTL 8 bit latches arranged as dual 32 bit wide Parallel-In-Parallel-Out (PIPO) shift registers, one for input and one for output, and two 256 bit wide ram I/O storage registers. Thirty-two bit wide words are written to the FIFO by presenting the data word and a strobe. As each new data word is written to the FIFO the previous data word is shifted over into the next 32 bit latch (See Figure 3.). When the eighth shift, or FIFO write, has occurred, all eight 32 bit words are transferred to the 256 bit input storage register. This 256 bit word is written to memory on the next available memory cycle. Eight more writes may now occur to the FIFO before another memory write cycle is needed. This scheme allows FIFO writes to occur eight times as fast as the memory cycle time.

The output configuration of the FIFO is similar to its input, as each 32 bit word is read out the next data word is shifted over to take its place. When the last word is read from the PIPO shift register several internal strobes occur that load new data from the 256 bit output storage register into the PIPO output shift register, and new data from the memory array is loaded into the output storage register on the next available memory cycle.

A unique feature of this FIFO design is that the larger the FIFO becomes the higher its data rate can be (within the limits of the AS TTL latches). The FIFO is being expanded to 16 megabytes which will increase the maximum data rate to over 100 megabytes/second. Doubling the internal word size to 512 bits allows a main memory read or write to occur half as often.

Memory Control

Memory cycle arbitration for read, write or refresh cycles is determined by a priority encoder. Writes have the highest priority with reads next and refresh with the lowest priority. While the current memory cycle is occurring, the decision is being made

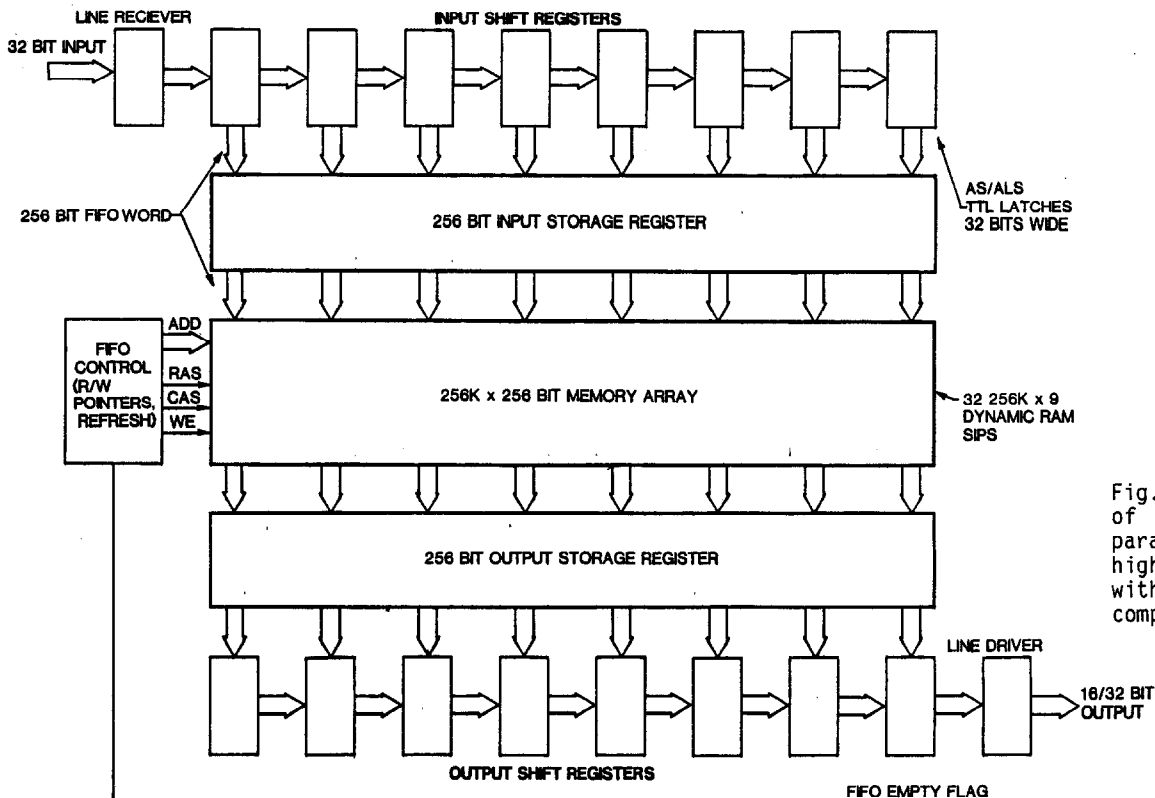


Fig.3 Block diagram of 8 megabyte FIFO, parallelism allows high overall speed with low speed components.




Fig.4 Reconstructed images of cosmic rays in a cerium doped glass target.

as to what the next cycle will be, and all memory addresses for read, write and refresh are being calculated. By the end of the current memory cycle, read, write and refresh memory addresses for the next cycle have been strobed into their respective latches. As the new memory cycle begins the address latch outputs are gated onto the address bus and the appropriate memory control signals (RAS, CAS or WE) are generated depending on the cycle type. When the FIFO is emptied of all data the refresh control logic is inhibited preventing memory refresh cycles and lowering power consumption.

This FIFO may be applied to any area where a very large high speed buffer is needed. Applications may include experimental data input buffers for high energy physics, medical imaging, high speed computer links, high speed data logging, ect.

Results

The only tests of the system for digitizing data have been with cosmic rays. A cosmic-ray telescope was set up to produce a crude trigger and gate the image intensifier and digitizing system. Software was written to reconstruct the image using compacted data read from the FIFO. Early tests were done with no compaction to verify the operation of the FIFO and image reconstruction software. Once these were fully checked, compactor operation was checked on test images and finally upon real time cosmic ray events.

The camera used was a Fairchild 3100F CCD with fiber-optic input plate, interlaced fields, and a resolution of 384 X 244 pixels/field (Since the final system will be non-interlaced we treat each field as a separate image.), and a pixel rate of 15 Mhz. After adjustment of the ADC, the camera was found to have 3 to 4 counts, (out of 64), of noise, apparently associated with the temperature or electrical noise in the environment. With this noise level in the camera, the digital threshold for the compactor was set to 3 counts and the video control section was programed to the proper field size for the CCD. With 384 by 244 pixels/frame the non-compactd CCD image would take 93,969 bytes to store.

Reconstructed Images

Figure 4 shows reconstructed images of cosmic-ray tracks. Fig. 4a required only 2432 bytes to store, for a space savings of 97%, tracks b, c, and d took 2784, 1728 and 2432 bytes respectively for savings of 97%, 98% and 97%.

The system was then used to examine the differences in various fiber-optic plate targets. We observed many tracks in several targets in an effort to determine the optimum target configuration. During these tests the compactor as well as the FIFO performed to our expectations, becoming powerful tools for the study of image data created by the scintillating glass targets. In one test, image reconstruction was halted and the system was allowed to accumulate cosmic-ray trigger data in the FIFO for many hours. We then read out these triggers in minutes looking for interesting events.

The FIFO has the capability to store at least 2000 frames of compacted data. This feature is needed during running of the actual experiment at Fermilab since the beam duty cycle is 20 seconds of beam followed by 40 seconds without beam. The FIFO stores data during the 20 second beam spill, and the computer transfers it to the optical disk during the 40 seconds between spills.

Optical disk data rates dictate the optimum size for the FIFO to be 16 megabytes. An upgrade of our present system is in progress to expand the FIFO to 16 megabytes and increase the ADC to 100 MHz.

[1] R. Ruchti, et al., "Scintillating Glass Fiber-optic Plate Detectors for Active Target and Tracking Applications in High Energy Physics Experiments," IEEE Transactions on Nuclear Science, Vol. NS-32, No. 1 (1984) 590-594.

[2] R. Ruchti, et al., These proceedings.