SLAC-PUB-1165
(EXPI)
December 1972

DESIGN OF A DATA ACQUISITION SYSTEM FOR LASS*

G. Armstrong, R. G. Friday, D. Hutchinson, K. Mauro, S. Shapiro, and S. H. Williams

Stanford Linear Accelerator Center
Stanford University, Stanford, California 94305

Introduction

The large aperture solenoid spectrometer (LASS, Fig. 1) now under construction at SLAC will put a strain on already existing data acquisition equipment to handle the projected 100 kiloword per second data rate. In order to meet this higher data rate and also provide a flexible data system to handle multiple users a new data acquisition system is being designed.

LASS is formed of two large magnets (Fig. 2). The first is a superconducting solenoid having a field region approximately 12 ft long with B=25 kG in the beam direction. The incident beam passes through a 1 meter hydrogen target at the upstream end of this solenoid. The secondary particles resulting from an interaction are trapped in helical orbits in the solenoid and their momenta measured by measuring the curvature of these helices. As high energy secondaries with small production angles have little curvature and cannot be measured well in the solenoid, a conventional dipole follows the solenoid to allow measurement of their momenta.

Thus a varied array of detectors within the solenoid, and fore and aft of the dipole, is required to reconstruct the particle trajectories of an event.

To measure wide angle tracks the target is surrounded by concentric cylinders of proportional chambers and spark chambers using capacitor-diode (C-D) readout. These are under design and production by our Cal-Tech collaborators.

Downstream of the target, at 75 cm intervals along the solenoid axis, are modular detectors, each composed of 2 gaps of 6'×6' spark chamber with capacitor diode (C-D) readout, plus 3 gaps of proportional chamber, 1'×1' centered on the beam axis, where particle density is highest. Proportional chambers are used to distinguish desired tracks from out-of-time background because their time resolution is 0.1 μsec, compared to the 0.5 μsec resolution of the spark chambers. It is not possible to use proportional chambers throughout because their complex readout costs about $7/wire (one wire per mm of chamber periphery) so the bulk of the solenoid is covered by C-D spark chambers which costs about $1/wire. Magnetostrictive (M-S) readout would be less expensive yet, but C-D readout must be used because of the high magnetic field in the solenoid region.

Between the solenoid and dipole there are 3 more proportional chamber-spark chamber modules of the same design as the solenoid detectors. The spark chambers may have magnetostrictive readout if it is found possible to shield the M-S wire adequately, but probably will also be C-D readout. Also placed between solenoid and dipole will be a 1.7 m × 2 m proportional chamber hodoscope being built by our Johns Hopkins collaborators. To keep costs down the wires will be spaced 4 mm apart and tied together in 1 inch groups per readout.

Following the dipole are four spark chamber modules using M-S readout. Two 5'×10' chambers are to be transferred from an existing spectrometer and two 7'×13' chambers will be built.

Also included in the post dipole detectors are two scintillation counter hodoscopes of 72 counters each, and a second Johns Hopkins built proportional chamber hodoscope 1 m × 2 m.

The final detector in the system is a large threshold Cerenkov counter containing eight photomultiplier tubes with an analog-to-digital converter on each tube.

In addition to these large detectors there are many smaller ones, beam particle position and direction proportional chambers plus various scintillation counters that must also be interrogated after each event.

Chamber Readout

C-D Chambers

The readout system for the capacitor-diode spark chambers (Fig. 3) is based upon that of Pizer.[1] Each wire has an individual circuit consisting of R_1, C_1, D_1, and R_D. A spark charges the wire side of C_1 positively, the other side being grounded by D_2. When the spark quenches, C_1 discharges quickly to +12 V through D_1 and D_3, then the discharge current is reduced as D_1 and D_3 become back biased and R_4 (1 MΩ) is the only remaining discharge path, so C_1 remains charged to 10 - 12 V during the readout time.

Readout is accomplished by clocking a single bit (READ) along a shift register chain.[2] Each shift register element switches a transistor into saturation which grounds the word bus, connecting 32 adjacent wire circuits. This grounds all 32 wires through D_1 and R_1 which will cause a negative pulse to appear on each of the data bus lines connected to a charged C_1. This circuit differs from that of Pizer in the addition of R_1, to protect D_1 from overvoltage and overcurrent, and R_D to help suppress ringing. Also, a transistor is used as a readout driver instead of a transformer.

The data bus and shift register chain is connected to a spark chamber control box which collects and formats the data before transmitting it to the PDP-11. The control box has its own clock and operates independently of the PDP-11 for most of the chamber readout. After the spark chamber is pulsed, the logic in the control box is held on "CLEAR" for 100 μsec while spark transients decay. Then the READ signal is clocked along the shift register chain, progressing down the side of the spark chamber at about 0.5 inch/μsec, until data is encountered. The data bus loads a 32 bit data register in the control box. When a nonzero word is loaded into this register the control gates off the clocking of the shift register chain. By reading out the data register serially and by using conventional logic and scalers the bit pattern in the data register is converted into three scaler readings.[3] Scaler 1 records the word number of the data (i.e., the number of clockings from the beginning of the spark chamber), scaler 2 the position of the first data bit within the word and scaler 3 the width of the spark. There is logic in the control to handle cases such as spark width overlapping the boundary between two words or more than one spark in the word. The data in these three scalers is formed into a fifteen bit word and transmitted to the PDP-11 via handshake mode. The control box then resumes clocking the shift register chain. The last shift register in the chain is connected to the control box to generate a "DONE" signal. All the ground side wires planes are read out by one control box; the shift register chains being connected in series. Similarly there is an identical system which reads out the high voltage side of the gaps. Read time is 1.2 msec, to

*Work supported by the U. S. Atomic Energy Commission.
(Presented at the IEEE 1972 Nuclear Science Symposium, Miami Beach, Florida, December 6-8, 1972)

scan the solenoid planes, plus 25 μsec (average) per spark times the number of sparks. For 120 sparks this is 1.2 ms + 3 ms = 4.2 ms.

It is in the test procedures that this system is most innovative. To avoid the space and cost problems of attaching one more resistor to each wire to put in a test charge, the reverse leakage properties of D_1 are utilized. A ten second negative pulse is put on each data bus in turn. Leakage through D_1 will charge C_1 positively. The readout sequence is started by the end of the test pulse. Readout failure of this bit in any word indicates an open D_1 or shorted C_1 or shorted wire. Readout of any other bit indicates a shorted D_1 for this bit. Testing for shorted D_1 can also be done by using a 1.0 μsec pulse on the data bus, charging C_1 through R_3 and D_4. If the reverse resistance of D_1 has decreased to 100 Ω or less (typically 200 MΩ) then C_1 will be charged by the short pulse.

Magnetostrictive Readout

In a previous paper[4] a system was described in detail which used a single counter and an inexpensive 20 MHz memory to store individual spark arrival times and the identification codes. A completed system is currently in use and has the capability of storing the time (position) data of 1028 sparks from 64 different magnetostrictive wands and in addition provides a real time CRT display of the stored spark information.

The system (Fig. 4) uses a set of multiplexed 5 MHz shift registers to form the memory element. This memory records the value of a single 20 MHz counter and a wand identification code whenever a spark signal is detected on one or more of the 64 inputs. Included in the system is a scheme for compacting the wand identification code, to reduce the size of each memory word, and a method for resolving the arrival and identification of coincident spark pulses.

Basically, the unit begins a readin cycle with a master event signal from the experiment. This signal resets all functions and starts the 20 MHz clock. As spark inputs arrive, the time counter and wand information are strobed into memory. The readin of data is ended by the time counter overflow signal which sets the unit into an advance mode to transfer data in memory to the end of the registers for output. When the data reaches the end of the shift register memory it is read into the computer. Finally the CRT display is refreshed by recirculating the data through the memory while waiting for the next master coincidence signal.

The data arrives at the direct memory access (DMA) in two consecutive words. The first word contains 15 bits of time and one bit of coincidence information followed by a second word containing bits of wand identification code. However, the DMA memory address register (Fig. 5) will automatically place only the 15 bits of time and one bit of coincidence information into computer memory in order to reduce the amount of recorded data. The computer will locate a 1 K page of core by presetting the page address register. This 1 K page in core is arranged into 64 wand tables containing 16 sparks per wand table. The 1 K page address of 10 bits is composed of the 6 most significant bits of wand identification address information and 4 least significant bits representing the number of sparks from that wand. A simple 64 × 4 bit scratch pad memory stores the number of sparks from each wand. The 6 bit wand address drives the scratch pad memory address lines and the scratch pad memory output forms the least significant 4 bits of the page address. Then the DMA using this concatanated address places the time data in the proper location within the 1 K page and the scratch pad output is incremented. The computer under program control reads the 64 word counts contained in the scratch pad memory and can perform further data compaction by moving the tables which don't contain 16 sparks together.

Data Acquisition Hardware

The large aperture solenoid spectrometer data acquisition equipment is designed to handle the projected 100 kiloword per second input data rate to the SLAC computer center. The overall computer arrangement is shown in Fig. 6. The data is collected by controllers located either in the control room or in the experimental building. These controllers pass their data to a DEC PDP 11/20 minicomputer which assembles and compacts the data for transmission via an IBM System 7 over a high speed coax data link to the S/370-360 triplex. Event data is also sent to an IBM 1800 at a much slower sampled rate for local monitoring and display. A control console containing a computer control panel and a keyboard, provides an operator interactive control over data flow, analysis, disk storage, logging and display during data taking runs.

PDP-11/20

In order to meet the specified data rate and also provide a flexible data system to handle multiple users the new data acquisition system utilizes the asynchronous unibus* of a PDP-11/20 (Fig. 7). The PDP-11/20 is required to provide the "front end" buffer of the experimental data to both the IBM 1800 and to the IBM System 7 terminal which drives the link to the SLAC computer center. The PDP-11/20 has been chosen because of its unique architecture which utilizes a unibus and a single instruction set that allows data to be transferred and/or modified between CPU, memory, and peripheral devices, with either the CPU or peripheral device having control of the bus.[5] Due to the asynchronous nature of the unibus it is compatible with devices that operate over a wide range of speeds. The asynchronous unibus is required by LASS because the experimental data from the various controllers arrives at different times.

The event data flows from the apparatus into buffer tables in the metal oxide semiconductor (MOS) memory and asynchronously out of the previous event tables into the IBM System 7 and up to the S/370-360 triplex. The PDP-11 central processor is not required to act in a direct way during the event transfer. However, at the event completion it swaps buffers and initiates control allowing the next event trigger.

Group G at SLAC has made a considerable design effort modifying CAMAC bins to provide a simplified method of building unibus interfaces.[6] In addition they have built a very inexpensive MOS memory system which we plan to utilize. Therefore, we purchased only 4 K of core memory which is the minimum amount that DEC sells with the computer.

The unibus will be extended out of the main frame and the external 16 K MOS memory constructed out of the popular 1103 dynamic random access memories. A series of direct memory access (DMA) channels handle all high speed data transfer. These include a link to an IBM 1800 as a local online sample analysis computer, a link to a remote IBM/370 series analysis computer via a System 7 terminal, and several DMA's to the experimental apparatus for data collection. The DMA's will be linked together in a simple daisy chain priority scheme with the DMA closest to the computer having priority and those downstream inhibited by the upstream DMA bus mastership. The DMA's are constructed on a modular basis using a double width CAMAC module and required approximately 60 IC's each (Fig. 8). The rest of the devices on the bus communicate with conventional techniques involving program-controlled interrupt type transfers.

The controllers are required to respond to signals from the associated DMA and the event trigger. In preparation

*Unibus is a trademark of the Digital Equipment Corporation.

for an event the PDP-11 program sets up in the DMA the memory address register followed by the word count register which sets the READY/SYNC flipflop to READY. After the DMA's are initialized and the PDP-11 is ready for the next event it sends a pulse which sets the PDP-11 "o.k." flipflop. This flipflop allows the fast logic associated with the experimental hodoscopes to be gated on. The fast electronics determines if there is an event of interest and issues an EVENT signal. This EVENT signal gates on the proportional chamber latches, fires the thyratrons, resets the PDP-11 "o.k." flipflop and sets the EVENT flipflop in each controller. The controller responds by presenting to the DMA 16 data bits and a SYNC pulse which resets the READY/SYNC flipflop demanding a DMA operation. The DMA requests a memory cycle and reads in the data and increments the memory address and word count registers and sets the READY/SYNC flipflop. This READY/SYNC process continues until either the word count overflows or the controller has presented all of its information, thereby resetting its associated EVENT flipflop and sending a DONE pulse to the done interrupt circuit. After receiving a DONE signal from each experimental data input DMA, a done interrupt is sent to the PDP-11 which then processes the event. Table I below lists these controllers and associated equipment. The asterisk indicates controllers with variable length records. Under program control the PDP-11 reads the individual word count registers.

TABLE I

DMA Event Requirements

		Words
1.1	Buffer Strobes	
	a. Beam Hodoscopes	2
	b. Beam Cerenkov	1
	c. Scintillation Counter Hodoscope A+B Hodoscope	10
	d. Magnet Veto	2
	e. Miscellaneous Counters	5
1.2	Proportional Chambers	
	a. Cylindrical Proportional Chambers	10
	b. (2) Prop. Hodoscopes Johns Hopkins-Fast Logic	10
	c. Beam Pro. Chambers	40
2.	40 Point ADC High Press Cerenkov	10
3.	a. Solenoid Chambers, Individual Wire	250*
	b. Cylindrical Spark Chamber	150*
4.	Proportional Plug Chambers	100*
5.	M.S. Wire Chambers Upstream Downstream	1000*
6.	DMA output link to IBM System 7 with status word control	1600*
7.	DMA output to 1800 with status word control not every event	1600*

With the above in mind, the flexibility and power of the scheme can be appreciated. For instance, any type of memory can be interfaced to the bus and can have cycle times up to the limit of the bus itself. The CPU need not be included in data transfers; therefore, data transfers to analysis computers can be direct and peripherals such as a live display can refresh directly out of memory. Data transfer priorities can be changed by plugging into a different DMA slot. Lastly, the system can be upgraded at will by simply changing CPU's as the technology changes such as incorporating a PDP-11/45 with floating point hardware on the same bus and utilizing a bus breaker technique.

IBM 1800

As shown earlier, part of the data acquisition and analysis system for LASS is an IBM 1800 computer. The 1800 was designed as a process control machine and has 32 K of 16 bit words and fixed point arithmetic instructions only. Peripherals include a line printer, two tape drives, disk drive, a card reader-punch, 2 storage CRTs and a console typewriter-keyboard. It will be connected by links to the PDP-11 and the IBM S/370 main computer.

The 1800 was made part of the LASS system for several reasons. One is that we have considerable experience (and applicable software) with it acquiring and monitoring data from spark chamber based experiments, including a link to another computer. The 1800 is located at the LASS site giving quick, easy access to monitor results. LASS is expected to begin operation before the S/370 is available, thus for a period, the 1800 will provide the only online monitoring while at the same time it is logging data.

Event analysis will be on a sampling basis whether or not the 1800 is being used for logging. In the logging case an event is selected from the input stream and monitoring analysis is begun as a background task. Acquisition of event data from the PDP-11 and logging on the 1800's tape drives are interrupt driven tasks. Analysis of the selected event thus proceeds at a rate determined by the number of cycles available to the background, not used by the critical logging process.

When the 1800 does not have logging responsibility, analysis will begin with the 1800 requesting data from the PDP-11 on completion of the previous event. Without the logging burden, more core and cycles will be available.

The purpose of monitoring analysis online to LASS is to ensure that the many data detection devices are functioning correctly. Event analysis consists of unpacking the data while generating statistics on the number of sparks contributed by each plane of all the spark chambers. Histograms are generated showing hodoscope firings for each scintillator and analog outputs of each big Cerenkov phototube. Spatial reconstruction of events is accomplished by track finding through the spark chamber subsystems individually. Track segments are linked together by extrapolation through the magnets. Since not every plane is required to have fired to find a track, track data can be used to measure the efficiency of elements of the spark chamber system. These efficiencies are measures of how well individual elements and subsystems of the spark chambers are working. Thus experimenters can watch performance of the primary and critical data detection device and react in an interactive manner with the experimental equipment.

Some monitoring features are not event related but are important to a collection of events. The current in many of the beam magnets is made available to the PDP-11 and thus to the 1800 through the link, via a programmable digital voltmeter which reads the shunt voltages. The 1800 checks that shunt voltages are within tolerances of specified values. Significant occurrences in the experiment, such as number of events, number of beam pulses, etc., are counted by fast scalers. These scaler values are available to the 1800 via the PDP-11. The scaler values can be used by the 1800's monitoring routines to indicate beam quality and composition. The 1800 can monitor system performance as a function of time by recording on its disk a set of numbers descriptive of the system's performance over the period of a "run" (perhaps an hour). Ten or twenty sets of "run descriptive" data can be saved and looked at by the experimenter at will.

Requests for CRT displays or printed information are made to the 1800 through mnemonic commands typed on its

keyboard, or through a set of thumbwheel and toggle switches. One thumbwheel specifies a particular CRT display, another specifies an argument to the display routine. Most displays have a corresponding line printer output for hard copy which can be obtained by turning on a particular toggle switch.

IBM System 7

The IBM System 7 and associated equipment function as a 100 KC data terminal between the PDP-11 experimental control computer and the S/370. The System 7 has 8 K of 16 bit 400 ns solid state memory and an ASR 33 teletype. Therefore most programs will be loaded via the link to the S/370. The System 7 also has 3 digital inputs and 3 digital outputs as the front end for the experimenter to use and a sensor based control adapter (SBCA) as a back end controlling the 1/2 mile bit serial coax link to the S/370. Internal to the System 7 the front end is serviced by a programmed interrupt scheme and the output SBCA by a DMA.

One of the 3 16 bit digital inputs is used for programmed interrupt control and the other two for digital input data. In order to maintain a data throughput rate of 100 KC 16 bit words, two data words are presented in pairs from the PDP-11 to the 2 digital inputs for each read interrupt, thereby reducing the interrupt response time by almost one half to approximately 18 μsec per word pair. At the end of the read interrupt routine a bit is changed in one of the digital outputs which is phase detected to form a sync pulse to the PDP-11 DMA. When the PDP-11 word count register overflows, it sends an end of record interrupt pulse to both computers. In addition, one digital output serves for status information and the other for a slow link to the 1800 for hard copy printer output and display.

370/360 Triplex

Actually the S/370-360 triplex consists of two new S/370 computers connected to the existing S/360 model 91 computer. The sensor based control unit (SBCU) which interfaces the System 7 link to a S/370 computer may be switched to either S/370 computer providing reliable 24-hour-a-day 7-day-a-week service. The S/370 computers will either be model 165's or 168's depending upon present negotiations with IBM.

Experimental data is received from the System 7 by a small spooling program in the S/370 which records the data on a 6400 BPI tape device. Information regarding the status of the spooling program is transmitted at low rate to the System 7 teletype or on to the 1800's line printer. The type of information sent will be run status, tape drive status, tape label information, event count, record count, I/O error count and possibly a limited amount of online analysis as core and CPU cycles are available.

A more sophisticated offline analysis written for the S/360-370 series computers exists and has been used in the past to monitor experimental data with a one hour turn-around time. It has been necessary to drive (by auto) the tape to computer center and wait or return for the printed output. But with the data stream passing through the S/370 in LASS, it will be possible to attach the offline program to the spooling program and analyze events directly in an on-line sense.

The offline program contains a more efficient algorithm for reconstructing spark chamber information with least-squares fitting routines allowing an improved monitor of the efficiency and resolution of the spark chambers. Of course much more core is available than in the 1800 so much more information about the spark chambers, proportional chambers, hodoscopes, Cerenkov counters, etc., can be recorded and be made available to the experimenter. The information is available online to the experimenter by means of the 2250 scope in the control room. The information can also be saved on a disk at the S/370 and retrieved by the Tektronics 604 graphics terminal after the analysis program has ended.

The core and cpu requirements of the analysis program are much greater than the spooling program so the use of it will be determined by the demands of the experiment for special attention to a device that may be failing and the demands of other S/370 users for core and cpu cycles.

Acknowledgements

We would like to thank Prof. D.W.G.S. Leith for his support and encouragement; to Dr. D. Porat and Dr. D. Dorfan for many helpful discussions and Mr. L. Birkwood for technical information regarding modified CAMAC bins.

References

1. H. I. Pizer, Nucl. Instr. and Methods 93, 249 (1971).
2. T. A. Nunamaker, Rev. Sci. Instr. 42, 1701 (1971).
3. R. Coombes et al., Nucl. Instr. and Methods 98, 319 (1972).
4. R. G. Friday and K. Mauro, "A Digitizing and Memory System for Wire Spark Chambers," IEEE Transactions on Nuclear Science, Vol. NS-19, No. 1 (February 1972).
5. PDP-11 Unibus Interface Manual, 2nd Edition, January 1971.
6. M. Schwartz and R. Flexer, Stanford Linear Accelerator Center, private communication.

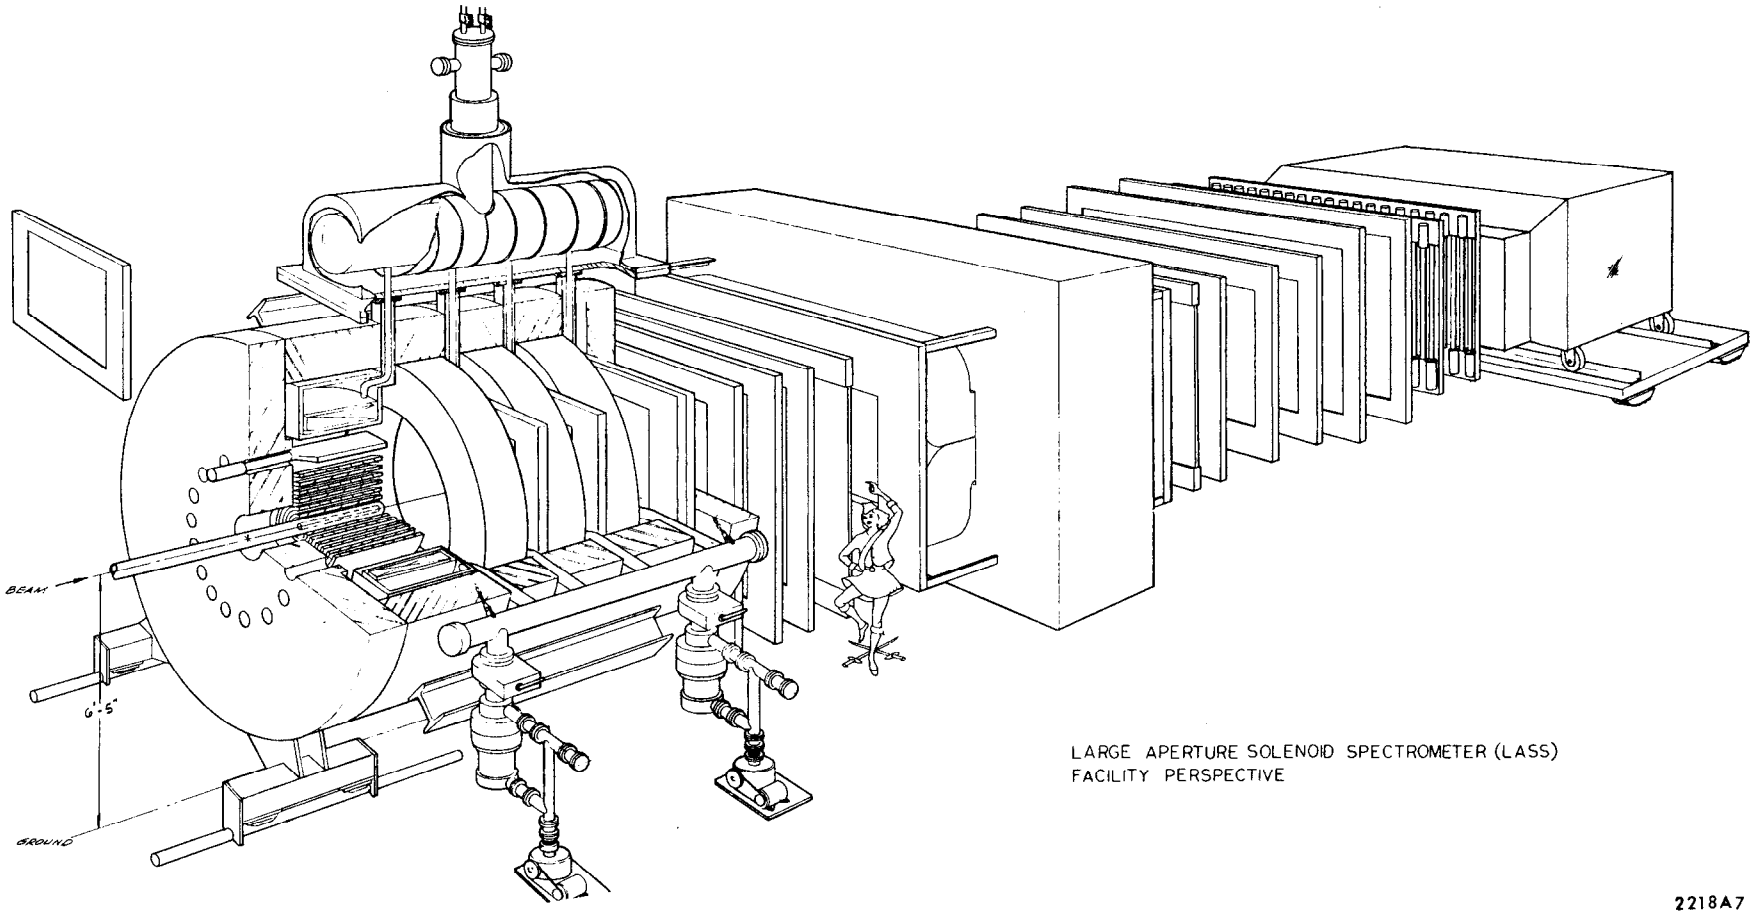

FIG. 1--Large aperture solenoid spectrometer (LASS) facility perspective.

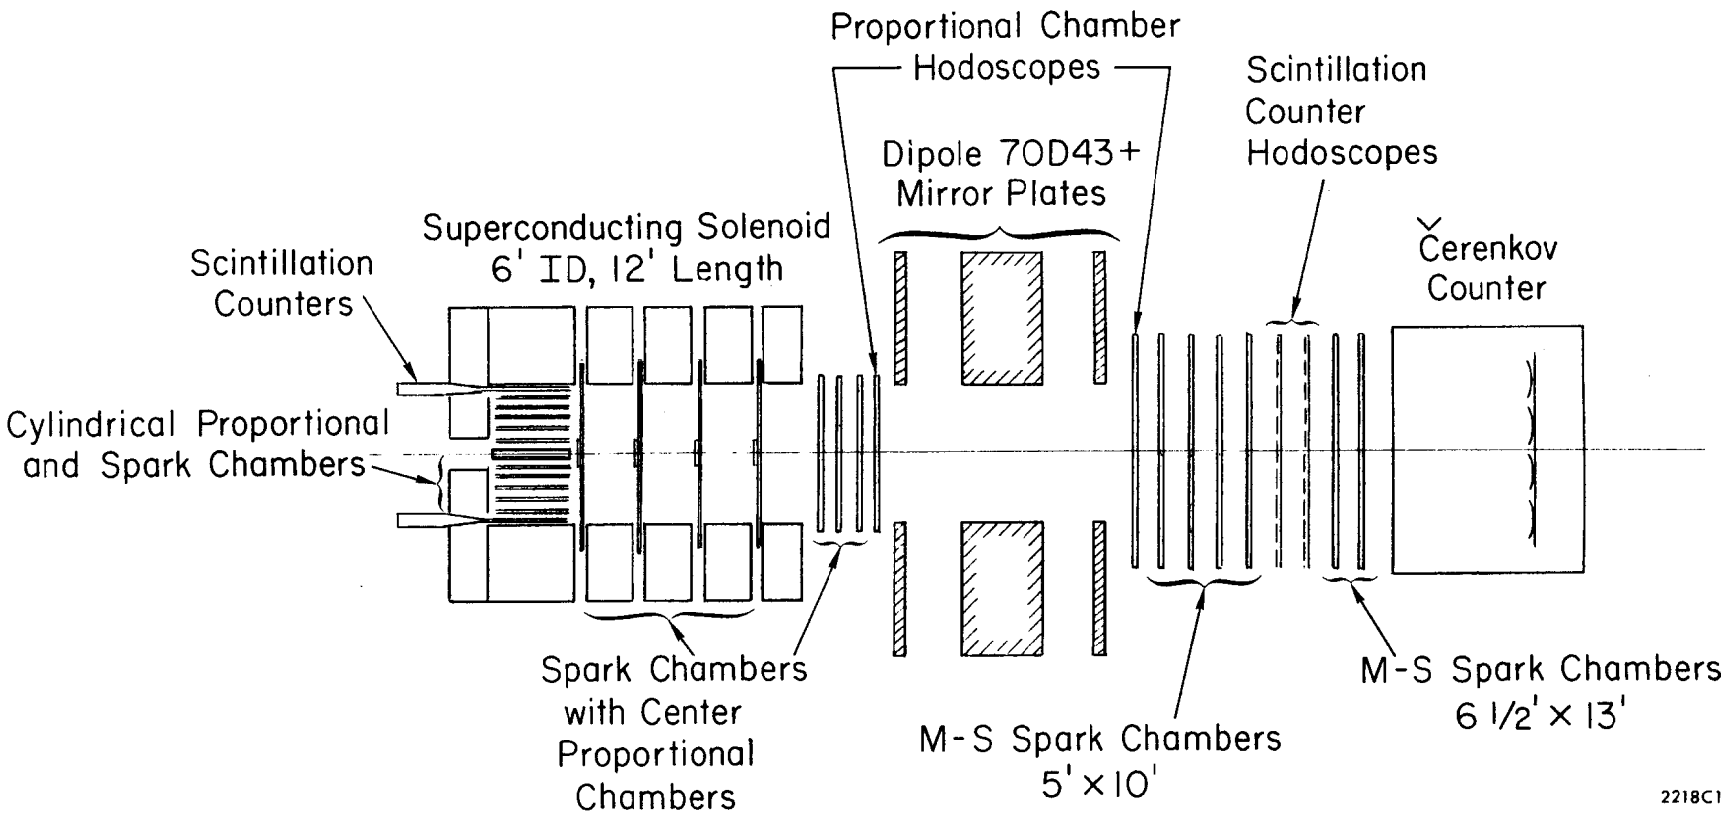

FIG. 2--LASS simplified plan view.

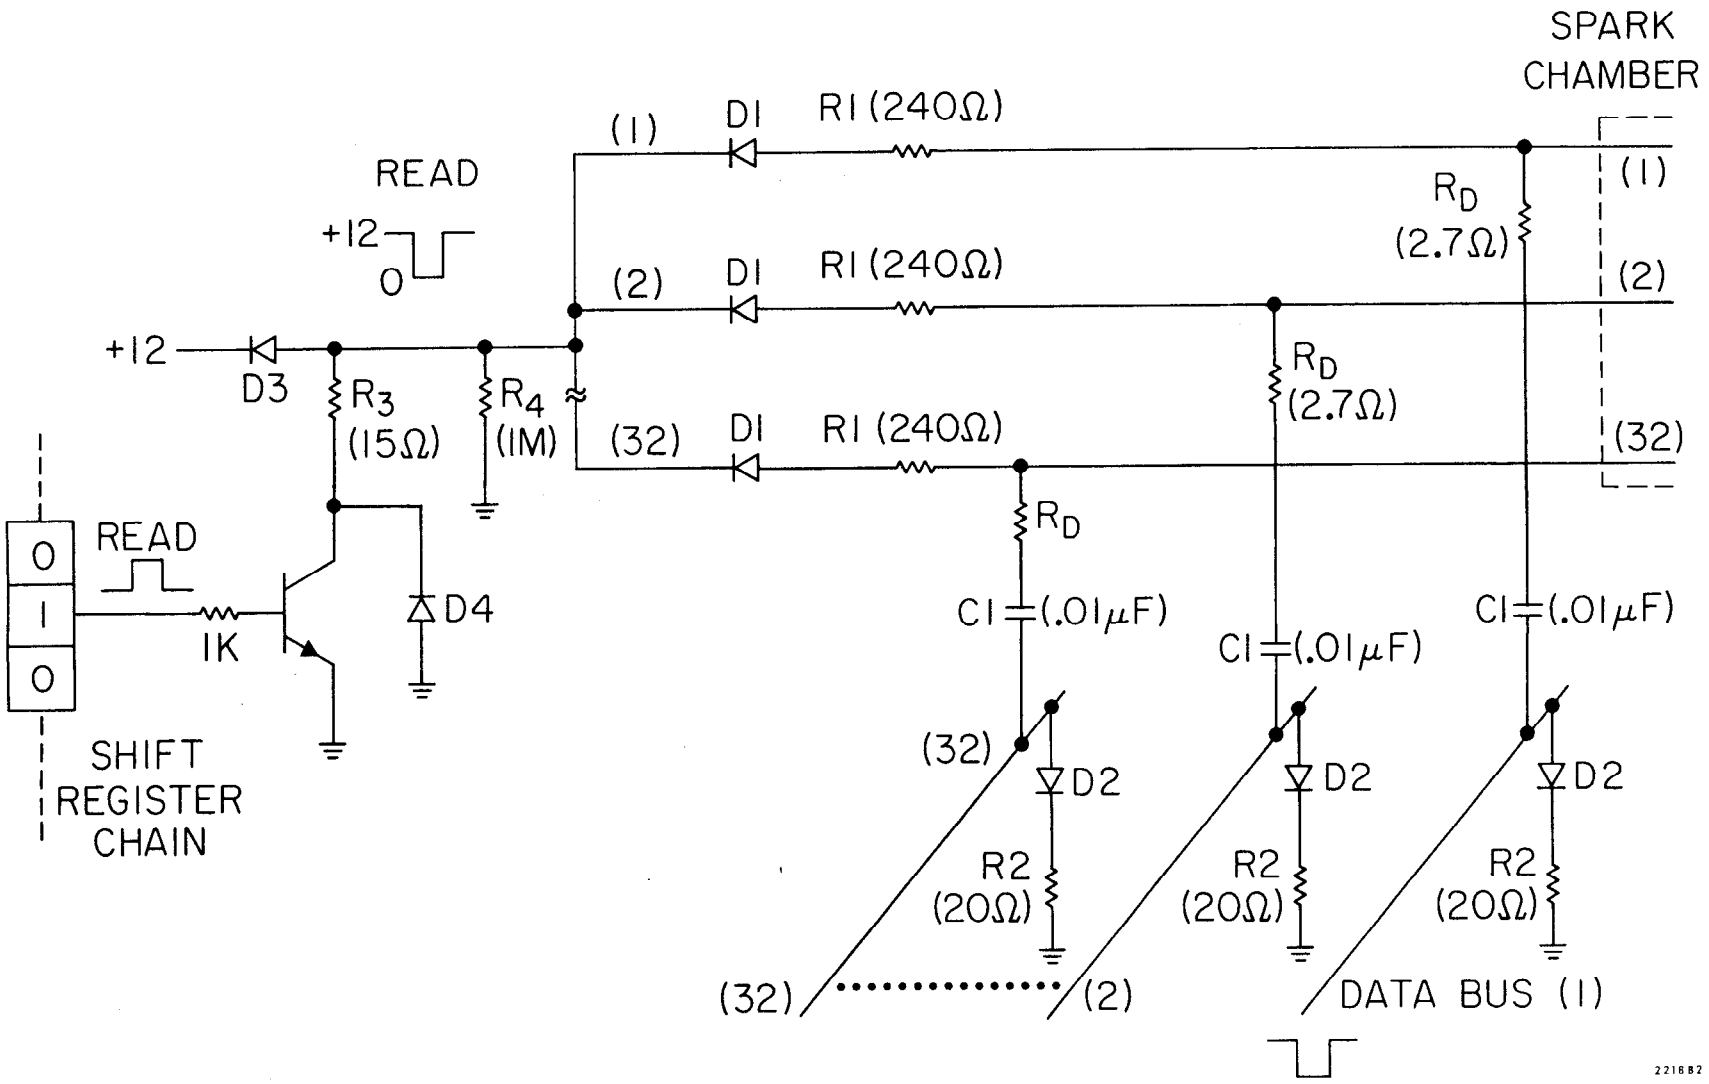

FIG. 3--Capacitor diode readout.

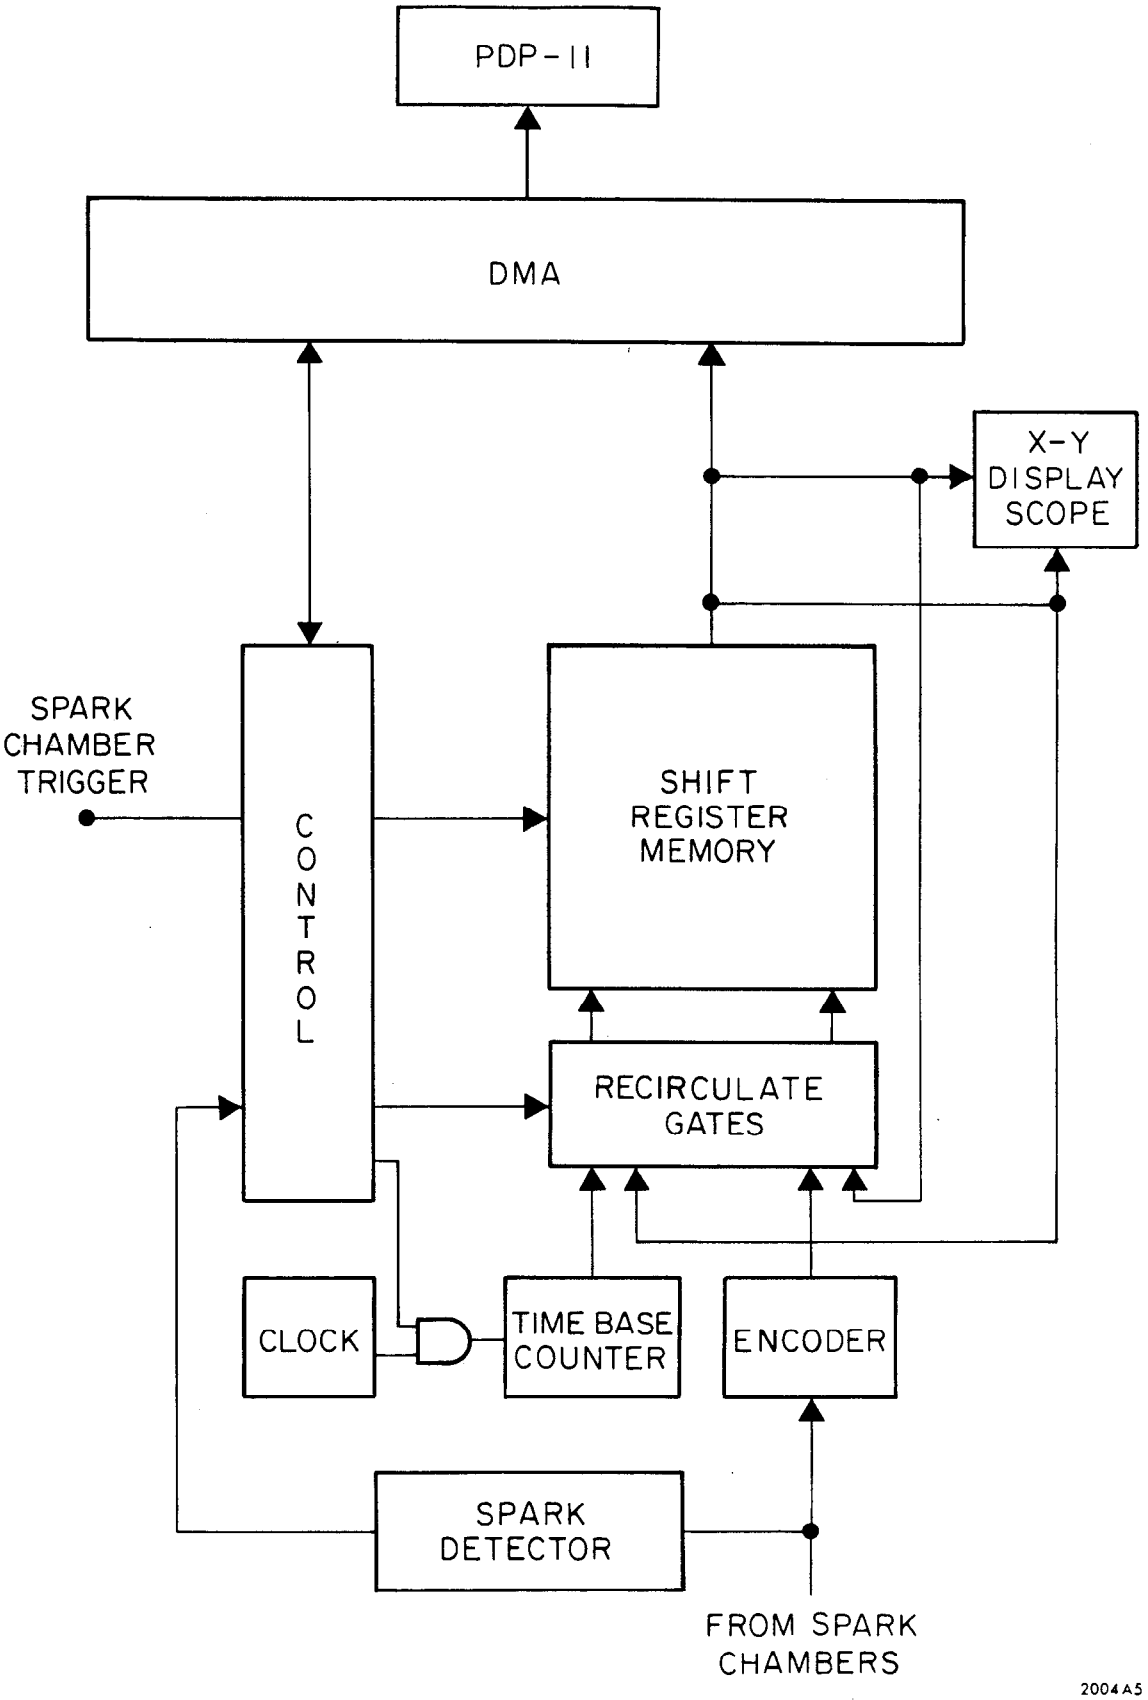

FIG. 4--Magnetostrictive readout block diagram.

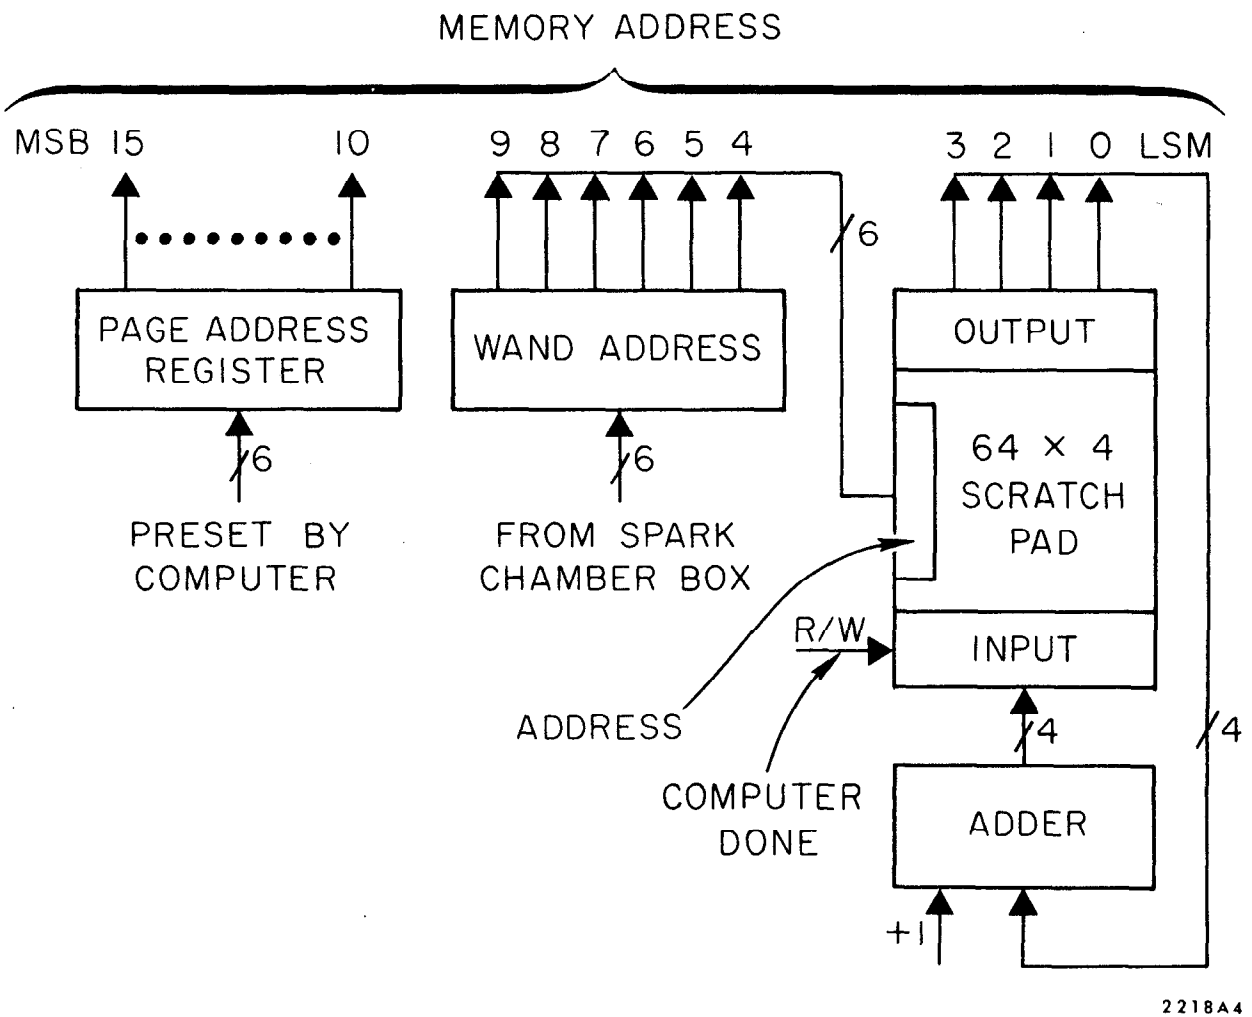

FIG. 5--Magnetostrictive address register.

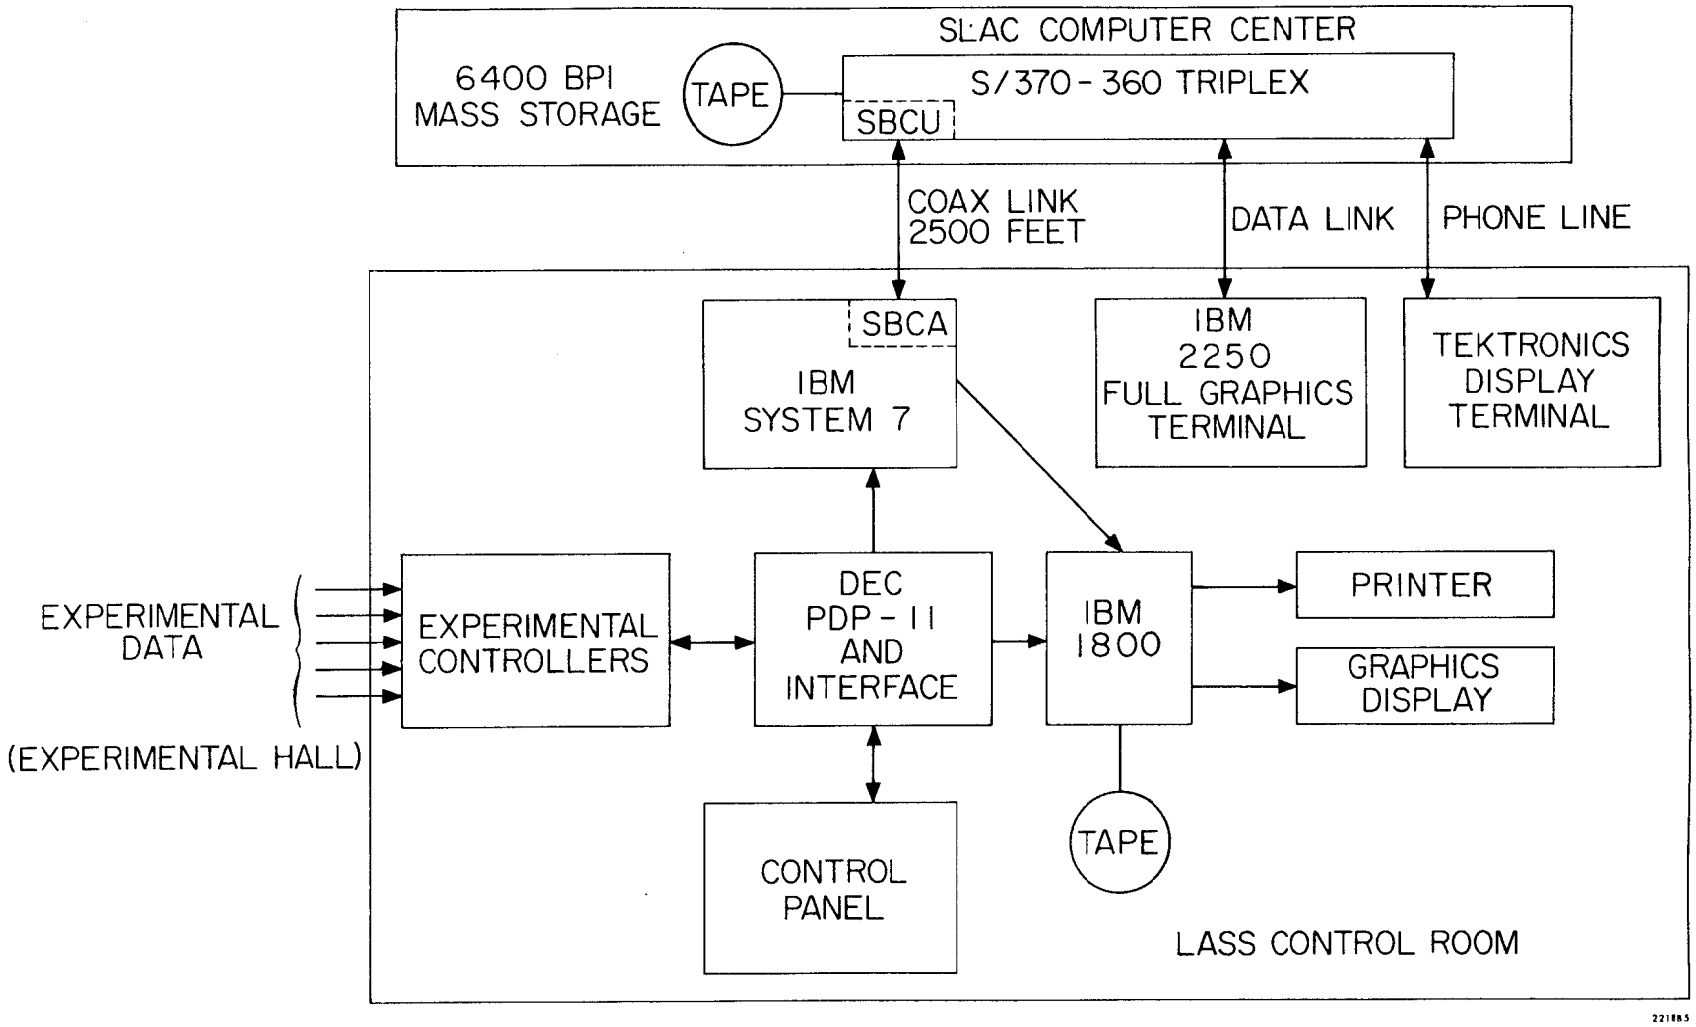

FIG. 6--Overall computer arrangement.

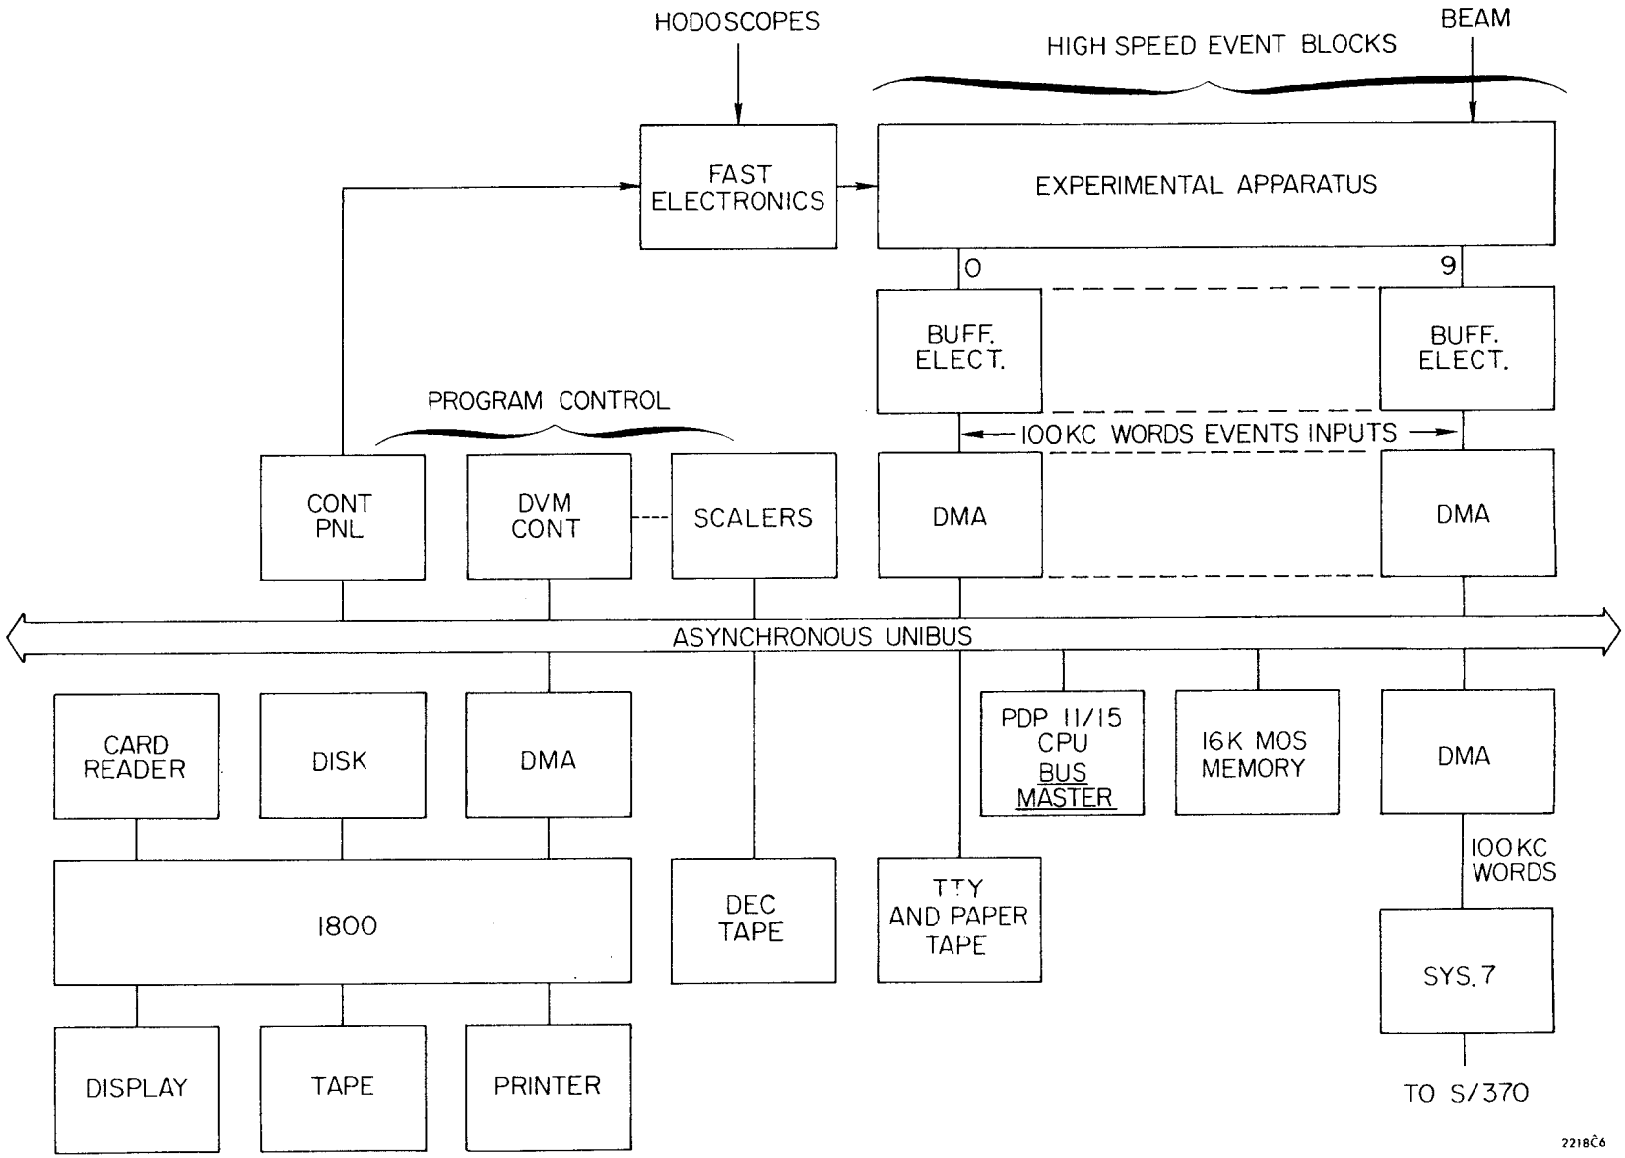

FIG. 7--Asynchronous unibus.

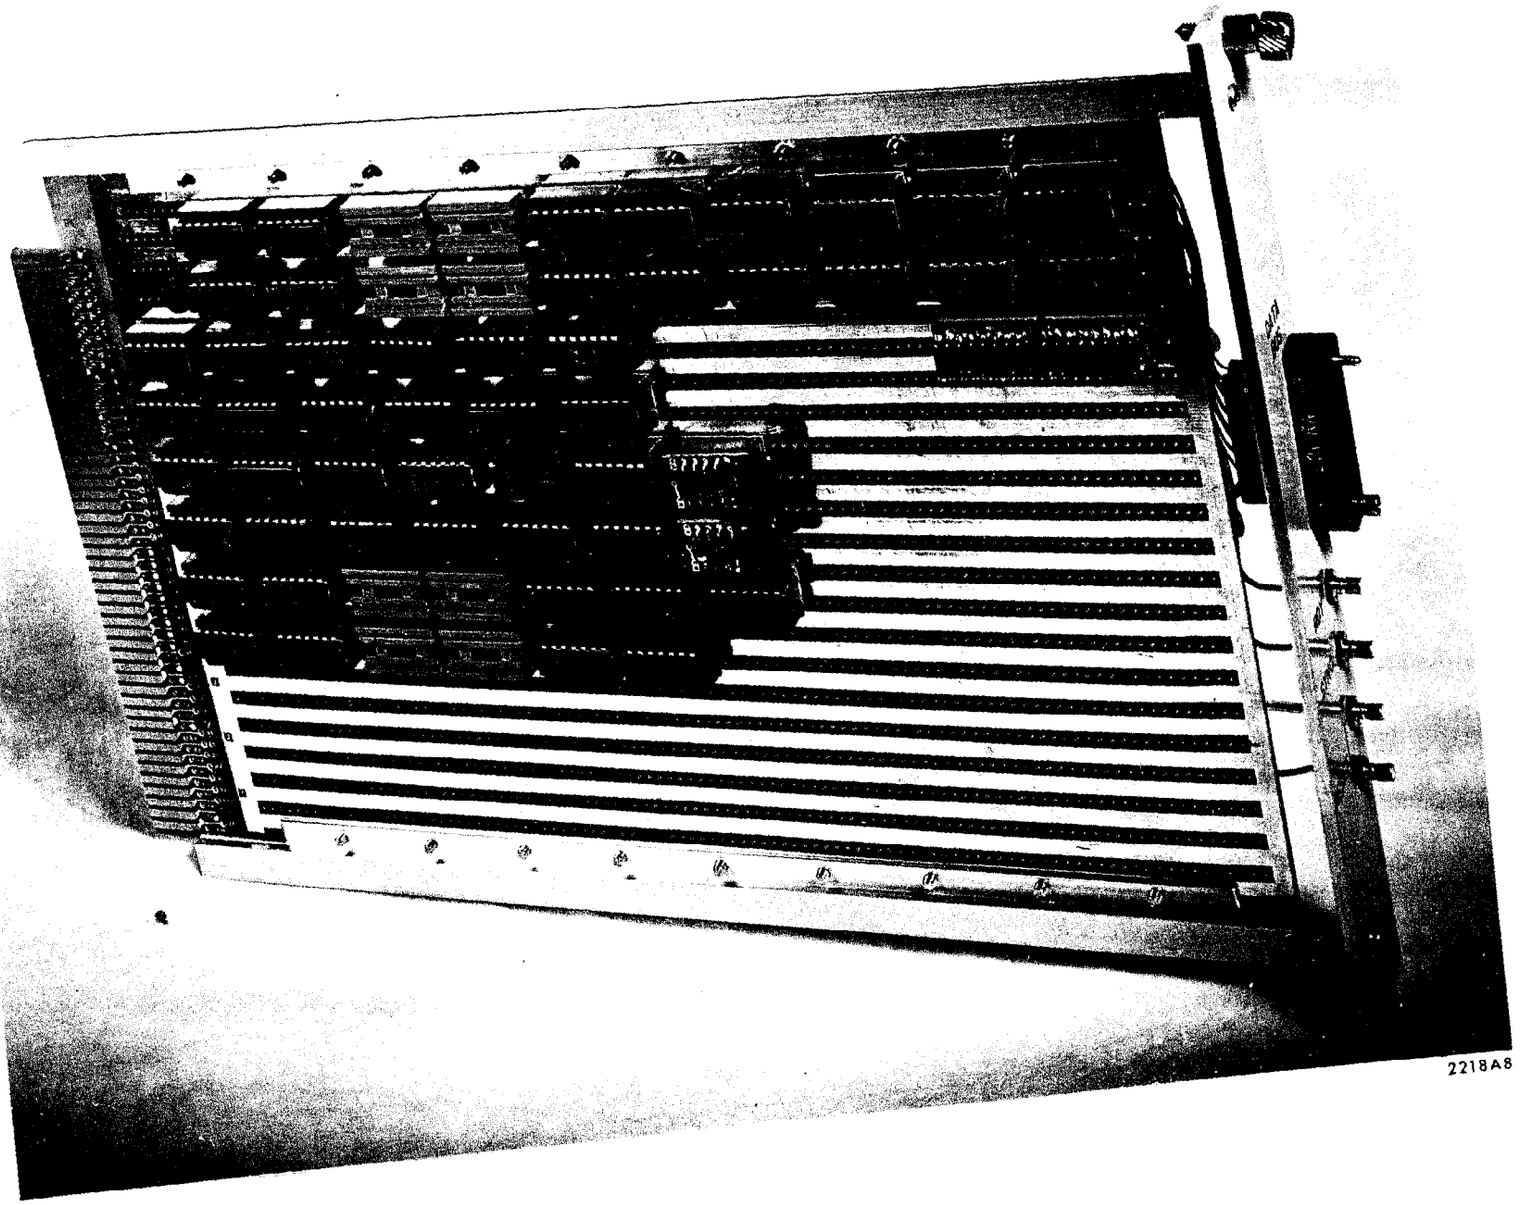

FIG. 8--PDP-11 direct memory address.