√

PREPRINT UCRL-78221

Conf-760006--15

Lawrence Livermore Laboratory

A MEGACHANNEL γ-γ COINCIDENCE SYSTEM USING A PDP-8/E COMPUTER AND MOVING-HEAD DISKS

W. D. Ruhter
D. C. Camp
L. G. Mann
J. B. Niday
P. D. Siemens

October 19, 1976

This paper was prepared for submittal to Proceedings of 1976 NSS/NPSS Symposia, October 20-22, 1976, New Orleans, Louisiana

MASTER

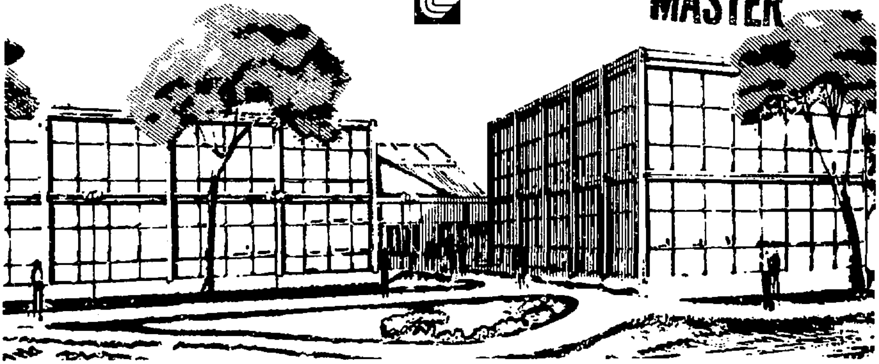

A MEGACHANNEL γ-γ COINCIDENCE SYSTEM USING A PDP-8/E COMPUTER AND MOVING-HEAD DISKS[*]

Wayne D. Ruhter[†]

and

David C. Camp,[**] Lloyd G. Mann,[**] James B. Niday,[**] and Philip D. Siemens[**]

ABSTRACT

A megachannel pulse-height analysis system using a PDP-8/E computer and two moving-head disk memories has been developed. The system has a storage capacity of 2^{20} memory locations, is capable of processing 1100 events/s, and provides on-line sorting and disk storage. An X- or Y-pulse-height spectrum in coincidence with one or several arbitrary pulse-height windows can be assembled in core for scope display and spectral analysis within 2 to 20 seconds. Reconstruction of a complete X- or Y-pulse-height spectrum requires about 3 minutes.

INTRODUCTION

With the advent of larger Ge(Li) detectors ten years ago it became possible to perform two-parameter γ-γ coincidence experiments using two Ge(Li) detectors. Gamma-gamma coincidence techniques used with NaI(Tl) detectors employed only 64 channels along each axis and, although marginally adequate for NaI based experiments, were totally inadequate for Ge(Li) based experiments. The order-of-magnitude improvement in energy resolution of the Ge(Li) detectors meant that at least 500 channels per axis were required or a minimum of 25×10^4 memory locations. Since the cost of ferrite core in this quantity was prohibitive ten years ago, alternative methods were developed.

Two-parameter systems were developed which allowed the simultaneous storage of approximately twenty 512-channel spectra.[1] Variable-width digital gates were set on selected gamma-ray transitions and/or neighboring Compton distributions. The disadvantage of this method is that one must pre-select the setting of the digital gates, and this may not be known a priori.

The buffer tape systems[2] were developed as an alternative method. In this technique, the entire information content of each event is listed sequentially in a core buffer, and then periodically transferred to a magnetic tape. With the buffer tape approach, there is, in principle, no limit to the resolution that may be retained for each parameter and no limit to the number of parameters.

Whatever method is used to obtain the coincidence data, the information normally desired is the pulse-height spectrum from one parameter (say X) in coincidence with an arbitrarily selected pulse height or pulse-height windows from the other parameter (Y). We will call this type of spectrum a slice of constant Y or a Y-slice. With the buffer-tape technique, the stored information must be sorted on large computers. This is both expensive and tedious, especially since only a limited number of gates per computer run of all coincidence tapes can be sorted. Since one good coincidence experiment may require 15, or even more, 2500-ft-tapes, it is clear that the data analysis will not be rapid.

The necessity for handling and changing tapes periodically can be eliminated by sorting and storing data on-line with the aid a computer and magnetic disk memory. The immediate availability of sorted data is advantageous in initially setting up an experiment and in monitoring the experiment. The principle advantage, however, is in the subsequent detailed examination of the data. Such systems have been developed with fixed-head/track disks[3] and moving-head disks.[4] These systems have been based on 18-bit/word and 24-bit/word computers, respectively, and are capable of processing 400 coincidence events/s with 2^{20} data locations.

We have developed a megachannel analyzer for on-line sorting of coincidence data with on-line data storage taking place on a moving-head disk. The basic system includes a 32 000, 12-bit/word, PDP-8/E computer; two moving-head disks, one with removable cartridge; two Analog to Digital Converters (ADC's); scope display with function box controls; line printer; and teletype. The system is shown diagrammatically in Fig. 1, while Fig. 2 is a photograph of the actual system. By coupling dynamic allocation of disk-track buffers with a compact arrangement of the more active portions of the data locations on disk, this system is capable of processing 1100 coincidence events/s.

DESCRIPTION

Disk Storage System

The magnetic-disk cartridges are divided into 406 tracks, each having 4096 words (12-bit), and each track is accessible to a movable read/write head. One disk provides 2^{20} locations (bins) in an XY array, each capable of storing 2^{18} events, while the other drive provides storage for programs and data files. Because the time required to update a track is 90 ms, it is important that the disk be used efficiently for both the data acquisition program and the data recovery and display program. The relationship between the disk addresses and the values of the parameters, X and Y in coincidence, is of primary importance in determining how efficiently the disk is used.

The coincidence data are stored in a 512 × 2048 array. However, 1024-channel resolution for the X parameter and 2048-channel resolution for the Y parameter is achieved by requiring that channel addresses, X_i and Y_i, generated by the ADC's for an X-Y coincidence event meet the criterion that $Y_i/2 + X_i$ be less than 1024. For those coincidence events having an X-pulse height greater than channel 511, the following algorithm is applied: $X_i' = 1023-X_i$; $Y_i' = 2047-Y_i$. This transforms the data from a 1024 × 2048 array to a 512 × 2048 array as shown in Fig. 3. Those events falling in region B of Fig. 2 are stored on the disk in an area corresponding to region A. With this approach all locations of the 512 × 2048 array are used rather than only half of a 1024 × 2048 array.

---

[*] Reference to a company or product name does not imply approval or recommendation of the product by the University of California or the U.S. Energy Research and Development Administration to the exclusion of others that may be suitable.

[†] Department of Applied Science, University of California, Livermore, California. Research supported by Fannie and John Hertz Foundation.

[**] Lawrence Livermore Laboratory, University of California, Livermore, California 94550

EB

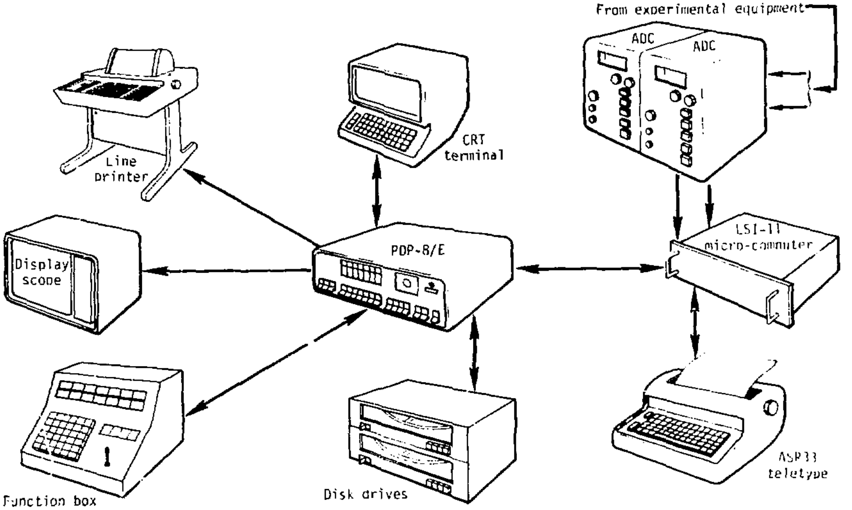

Figure 1. Diagram of computer system components. The system as shown allows two-parameter data acquisition.

To insure that all real X-Y coincidence events meet the above criterion, the Q-value of the nucleus being studied must fall within the 1024-channel resolution of the X detector and the 2048-channel resolution of the Y detector. This is accomplished by adjusting the amplifier gains of both detectors accordingly.

To avoid the word-and-a-half manipulation implied by an 18-bit data capacity per XY location and, at the same time, improve the data acquisition rate, each coincidence event location is separated into two parts. The least significant 12-bits of a data location reside on one track while the remaining most significant 6-bits reside on another track. The most significant part of a data location is involved in an update every 4096 events for that location. In this way only 256 tracks are needed to store the least significant 12-bits of 2^{20} locations in the array. The overflow events are added to a buffer which, when full, is emptied onto the 128 tracks used for the most significant part of each data location. This compact arrangement of the more active portions of the data locations results in less time being lost due to disk-head motion, consequently improving the data acquisition rate.

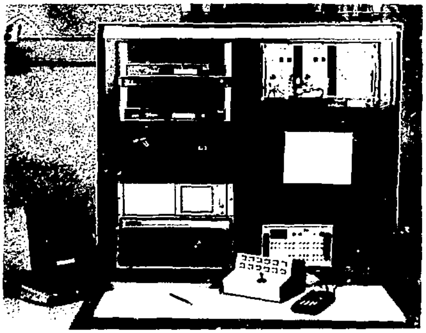

Figure 2. Photograph of computer system diagrammed in Fig. 1.

Once the channel addresses, X_i' and Y_i', have been determined from an XY coincidence event, it is necessary to convert this information to a disk track number, P_i, and a location in that track, Q_i. One approach would be to place the 2048-channel Y-spectrum that is in coincidence with $X_i = 0$ (i.e., the $X = 0$ slice) on the first half of track number 0, the 2048-channel $X_i = 1$ slice on the second half of track number 0, etc. Each track would then store two complete X slices. However, this disk format has a major drawback, namely, the time required to read a Y slice into core. To assemble a one-channel-wide Y-slice into core would require reading all 384 tracks or about 20 seconds. This becomes excessive when one considers the problems of not only assembling a spectrum in coincidence with a window that is several channels wide but also the time required for processing the background slices that must be subtracted from it. The solution to these problems is shown schematically in Fig. 4. With this scheme the data are partially randomized, while still allowing for reasonable data-recovery times in both directions.

The tag bit shown in Fig. 4 is used to indicate whether a coincidence is a true or chance event. If the event is a chance coincidence, then the appropriate location on the disk is decremented.

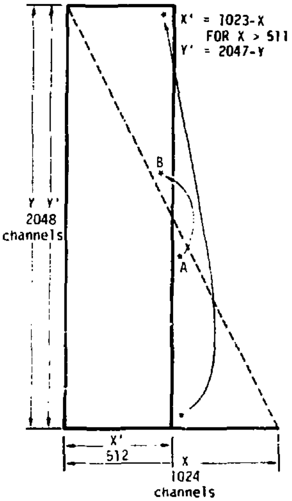

Figure 3. Diagram showing where coincidence events with X-pulse heights greater than 511 are stored in a 512 by 2048 array.

List Mode Processor

A method proposed by Gonidec[5] and used by Sinclair *et al.*[6] for "list mode processing" the coincident events from the ADC's is used with some modifications. The processing system uses the interrupt feature of the PDP-8/E along with a software priority-level scheme.

Since data transfers to and from the disk occur through a 'cycle-stealing' direct memory access mode, other computer operations can still take place at high efficiency. An entire track is updated before moving the head to the next track; however, only half a track is transferred at a time. This allows the first half of a track to be updated while the second half is being read in, etc. The completion of reading or writing one half-track by the disk causes the computer to be interrupted and the necessary commands for the next disk operation are issued within 33 μs. This is a short enough time to allow the disk to begin operations on the next half of the track without losing a revolution. The disk interrupt also requests the processor with the highest priority to assemble the instructions for reading or writing the next half-track and to update the half-track currently in the disk buffer.

Two lists, of 128 data words each, are cyclically opened to the ADC's via an LSI-11 microcomputer. The LSI-11 is used to determine whether the digitized events from the ADC's, X_i and Y_i, meet the criterion that $Y_i/2 + X_i$ is less than 1024, and, if they do, to transform them to X_i' and Y_i'. Once this is completed, the LSI-11 interrupts the PDP-8 and passes the two data words, X_i' and Y_i', to the PDP-8. The PDP-8 converts these two words to a half-track address, P_i, and an address on that half-track, Q_i, and stores the PQ pair in one of the lists. When one list is filled, the other is immediately opened to the ADC's while the first is emptied by the processor that is one level lower in priority than the disk processor.

At this processor level, the performance of two distinct functions is required. First, the incoming data must be organized into lists, each corresponding to a half-track on the disk, for efficient disposition. Second, the size of these lists must be variable to allow the data belonging to track T to be saved until the disk returns to track T. The time to update one track is 90 ms. Consequently, about 23 s are required to completely update the 256 tracks which contain the least significant portion of each data location. To satisfy these two requirements, two arrays are defined: a "fixed memory" for the first requirement, and a "free memory" for the second. Each of the arrays is composed of groups of eight words or "octets."

The fixed memory occupies 4096 words of core. Each of the octets in the fixed memory consists of 2 control words and 6 data words. The free memory occupies 12 288 words of core. Each of the octets in the free memory contains one control word and 7 data words.

The first word of a data pair in the ADC list determines the octet in the fixed memory where the other word of the pair is to be stored. If the first control word of the fixed octet is a counter, then the existence of an empty data word in the fixed octet is guaranteed. When the counter reaches 6, an empty octet in the free memory is found, and the first control word of the fixed octet becomes a pointer to the location of the new free octet. The second control word in the fixed octet points to the free octet that was last filled. The control word in this free octet points to the prior-filled free octet. In this way the octets in the fixed memory are assigned to consecutive half-tracks on the disk, while the free memory provides additional storage for each half-track by linking together a variable number of octets by pointers. All of the free octets are linked together initially. When one is emptied, it points to the free octet previously emptied. The octet just emptied becomes the first-available empty octet (see Fig. 5).

Overflow of the least significant part of a data location will occur when the number of counts in that location exceeds 4095. (Underflows can also occur when a chance coincidence is subtracted from a data location containing zero events.) When an overflow/underflow occurs, two words are stored in a list of 256 words along with a tag bit to indicate whether it was an overflow or underflow. These two words give the half-track number and the address on that half-track where the overflow/underflow occurred. Once the list is full, the highest priority processor is switched from updating the least significant tracks to updating the 128 tracks used for the most significant part of each data location. When the list is empty, the processor resumes updating the lower 256 tracks as described above.

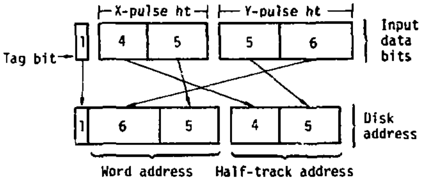

Figure 4. Disk format for semirandom data storage and optimum data readout.

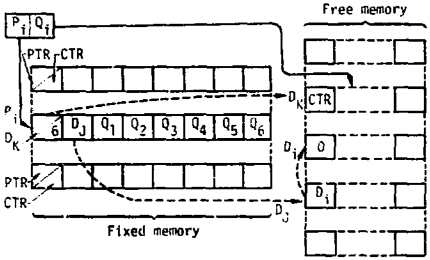

Figure 5. A schematic showing organization of fixed and free memories. The D_J points to the octet that was last filled while D_K points to the octet that is currently being filled. The 0 in the D_i free octet indicates that it is the last in the chain. The first control word in a fixed octet is used either as a pointer (PTR) or as a counter (CTR).

Besides the two processor levels which control the flow of data from the ADC's to the disk, there are lower priority levels which process the interrupts from the teletype and function box. The main purpose of the function box is to provide interaction with the scope display. The scope display program operates at the "background" level of the computer. While the data are being acquired, a 1024-channel spectrum is displayed showing those events observed in the Y-detector that were in coincidence with events in the X-detector. Either half of the 2048-channel Y-spectrum may be displayed by pressing a button on the function box.

The rate at which data can be acquired is determined by the size of the free memory and the speed at which data can be stored on the disk. If the 16 000-core buffer space had been divided into 384 spaces of 42-word buffers, then the system could have processed 466 events/s (42/0.09 s). With the dynamic allocation of track buffers this rate would have increased to 688 events/s (62/0.09 s). By dividing the data locations into two parts, the number of buffers needed was reduced to 256. consequently increasing the buffer size. This gives the system the capability of processing 1155 events/s (104/0.09 s), disregarding the time spent for overflows. If we assume 90 ms to process an overflow/underflow, the maximum rate at which the system can process is 1127 events/s (104/[0.09 s + (104/4096) 0.09s]).

Data Recovery

Once data acquisition is stopped, a separate program must be called in from the disk on the upper drive for examination of the coincidence data. This program can assemble an X(Y) spectrum in core that is in coincidence with up to four Y(X) spectral regions (windows). These spectral regions may be added or subtracted, thus allowing any particular gamma-ray coincidence spectrum to be examined with the adjacent Compton background region(s) subtracted if desired.

The amount of time required to assemble a spectrum in core depends on the number of windows and the channels per window that are requested. Typically, to assemble a spectrum in core requires several seconds. A full X(512) or Y(2048) window requires about three minutes. Rapid recovery of the data in a meaningful format is highly advantageous not only in monitoring an experiment but in subsequent data examinations by scope display or other read-out peripherals. Auxiliary programs which perform desired peak integrations, centroid calculations, or other specialized jobs can be called in from the disk on the upper drive.

SUMMARY

A real-time megachannel analyzer assembled from a 32 000, 12-bit/word computer and two moving-head disks has been described. By efficient use of the disk and by dynamic allocation of track storage, the system allows for a considerable input-data rate with good energy resolution. Spectral data are immediately available for detailed examination at any time. The system is flexible, and increases in core size or disk storage space can easily be made. Such increases would allow increased input-data rates or larger data arrays.

ACKNOWLEDGMENT

This work was performed under the auspices of the U.S. Energy Research and Development Administration, under contract No. W-7405-Eng-48.

REFERENCES

1. N. S. Kendrick, Jr. and D. A. McClure, "Multiparameter Pulse-Height Analysis with a PDP-8/E Computer", Nucl. Instr. and Meth. 121, 573-580 (1974).
2. D. A. McClure, N. S. Kendrick, and J. W. Lewis III, "A Three-Parameter Magnetic Tape Recording System", Nucl. Instr. and Meth. 103, 453-460 (Sept., 1972).
3. J. B. Niday and L. G. Mann, "Megachannel Pulse Height Analysis" in Proc. Intern. Conf. Radioactivity Nucl. Phys. (Gordon and Breach, New York, 1972) p. 313-322.
4. R. A. Schrack, H. T. Eaton II, and D. Green, "A Modular Minicomputer Multiparameter Data Gathering and Virtual Operating System for the NBS Neutron Standards Program", in Proc. Nuclear Cross Sections and Technology Conference (NBS Special Publication 425, Vol. 1, 1975) p. 97-98.
5. J. P. Gonidec, "A 512 000 Channel Bidimensional Analyzer Using an IBM 1800 Computer", Nucl. Instr. and Meth. 88, 125-135 (1970).
6. J. W. D. Sinclair, J. W. Smith, C. M. Rosza, and S. L. Blatt, "A Real-Time 448 000-Channel Multiparameter Pulse-Height Analysis System", Nucl. Instr. and Meth. 111, 61-65 (1973).