J. F. Mollenauer

Growth of a Laboratory Computer System for Nuclear Physics*

Abstract: A computer system may typically be expected to progress through a cycle terminating in overloading. The experience with an early system at the author's nuclear physics laboratory serves as an example. The original computer and a similar machine later installed with it are now overloaded and a new system is under construction. The success of the interactive data analysis on the original system has made it desirable to enhance the display and light pen facilities while reducing the computer time involved in generating the displays. The use of a data storage/display disc effectively provides off-line displays but requires more manipulation in data acquisition. The solution is found in the large number of processing units economically feasible with third-generation equipment. Two linked computers will perform data acquisition and analysis, the smaller performing data acquisition under the control of the larger, which will run a fairly simple time-sharing system. Together with several I/O processors, this hierarchy of processors will provide ease of program development and a very high degree of computational power and data acquisition capability.

Introduction

In the last few years the computer has assumed a position as the most prestigious of instruments in the research laboratory. Like other new instrumentation, a computer opens new approaches for its users, then is used routinely, and is finally outgrown by demands it has itself stimulated. However, the cost and time required for implementing a computer system greatly exceed most other instrumentation. Particularly when a computer is specially interfaced to experimental equipment, replacing it can be painful.

At the Bell-Rutgers nuclear physics laboratory we have been using one of the earliest laboratory computer systems in nuclear physics. By now this system is obsolescent and its replacement is currently being installed. Hopefully our experience in the usage of this system as reflected in its redesign will be useful in the implementation of other laboratory computer systems.

The programming barrier

The introduction of a computer has had two major effects on our research effort. On the positive side, it has made possible more complex experiments and more effective data analysis than would have otherwise been possible. (Examples of these uses will be given later.) On the other hand, the very flexibility of a computer raises a substantial barrier to the setting up of a research project. One may learn to use a wired-program multichannel analyzer in half an hour; to attain the same proficiency in programming and using a computer takes months.

Most of the flexibility of a computer system is thus unavailable on a practical basis for trying out new ideas in data acquisition or analysis. The use of FORTRAN with a card reader and a line printer gives our system an advantage over many small computers, but the programming barrier is still far too high. The only visible solution to this problem is a console-interactive editing and debugging facility with good secondary storage capabilities. The cost of such a computer system dictates that it be time-shared among a number of users. Adding to the well-documented difficulties of implementing such a time-sharing system are the requirements that it must also handle several independent real-time data acquisition tasks and provide displays for monitoring experiments and data analysis. The problem is made tractable only by the advent of high performance secondary storage devices and by the fact that the number of time sharing users will be small, of the order of four or five.

Usage of the present system

While the programming barrier has retarded implementation of many applications, the present system has been quite successful in many ways. It is now used for all experiments run on the laboratory's tandem Van de Graaff accelerator. Many of these applications could not have

The author is at Bell Telephone Laboratories, Incorporated, Murray Hill, New Jersey.

* This work was supported in part by the National Science Foundation.

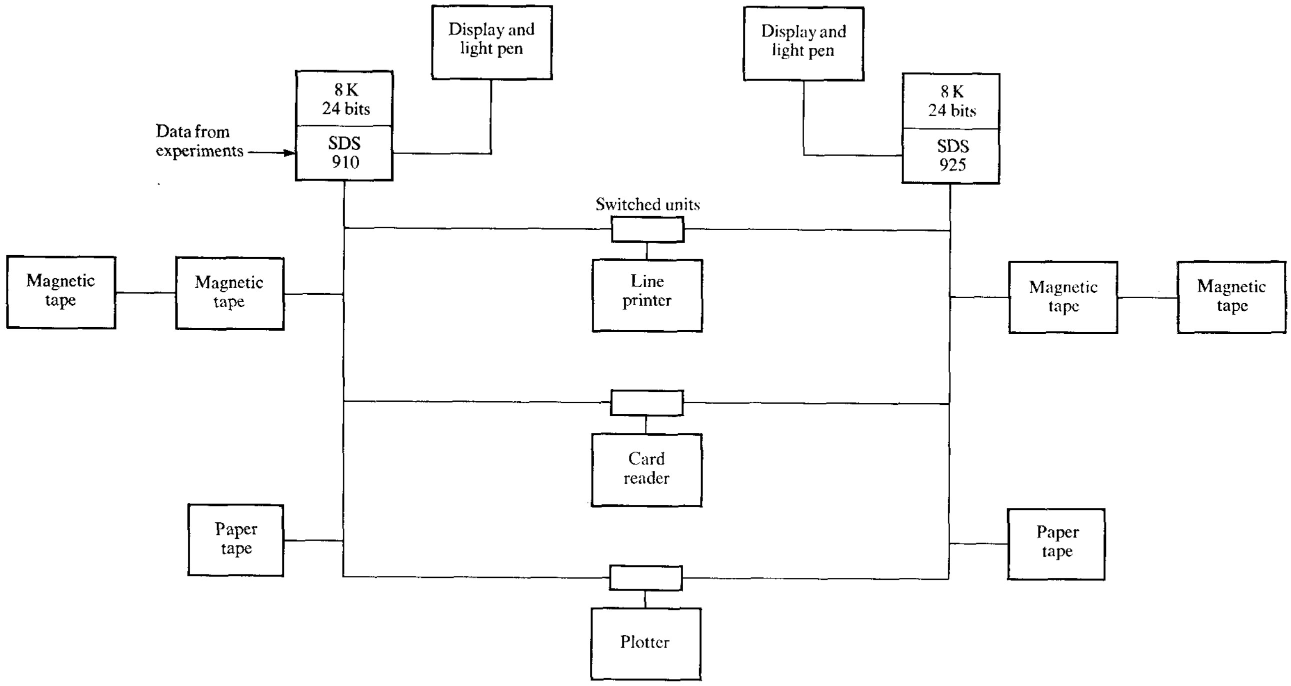

Figure 1 Schematic diagram of the present system with an SDS 910 and 925, showing switched peripherals. Initially the system included only the 910; the 925 was added later.

been attempted without a local computer and indeed were not foreseen when the system was first proposed. While these methods are by now familiar to nuclear physicists, they represented substantial advances when first adopted.

On-line data simulation

After the installation of our computer system in 1962[1] (Fig. 1) the first hint that the computer might revise our ways of running experiments came from the discovery that the display system could display quantities other than data from the experiment. In experiments dealing with the breakup of energetic nuclei into three particles, normally the energies of two of them are observed. When these two energies are observed at specific angles, one degree of freedom remains in the system—events from a given reaction are constrained to fall along a curve or "kinematic line" on an energy-vs-energy plot. Amplifier nonlinearities, calibration errors, and many other experimental faults can cause displacement of the curve from its expected position. The curves must be calculated and checked against the data to verify the experiment. While not difficult, calculating and plotting them was tedious until we found that they could be displayed by the computer in the same coordinate system as the data themselves. Checking the experiment became a simple matter of running the calculation and photographing the display for reference against the incoming data. This could not have been done without the computer. A great deal of valuable accelerator time was saved by detecting faults at the beginning rather than at the end of the experiments.

In addition to indicating the locus of a particular reaction in the energy-energy plane, this method has been extended to predict the yield of the reaction along the line. The effects of detector aperture and resolution, target thickness, and counting statistics have been incorporated as well. By comparing displayed predictions of various models with the experimental results, the conclusions have been drawn from many experiments without the necessity of working up the data any further. For greatest accuracy, conventional analysis methods have been used afterward, but "eyeball" adjustment of theoretical parameters has often given results good to a few percent. An example of a theoretical fit to a two-dimensional spectrum by D. P. Boyd[2] is given in Fig. 2.

Another advantage of the simulation method is that it provides an excellent way to determine the best conditions under which to operate the experiment. For example, widening the detector aperture improves the data collection rate but broadens the kinematic line, possibly washing out needed detail. By using Monte Carlo methods to average over the detector aperture and target thickness, the necessarily imperfect results of the experiment are predicted very closely. Thus the experimental parameters can be optimized before the experiment is run.

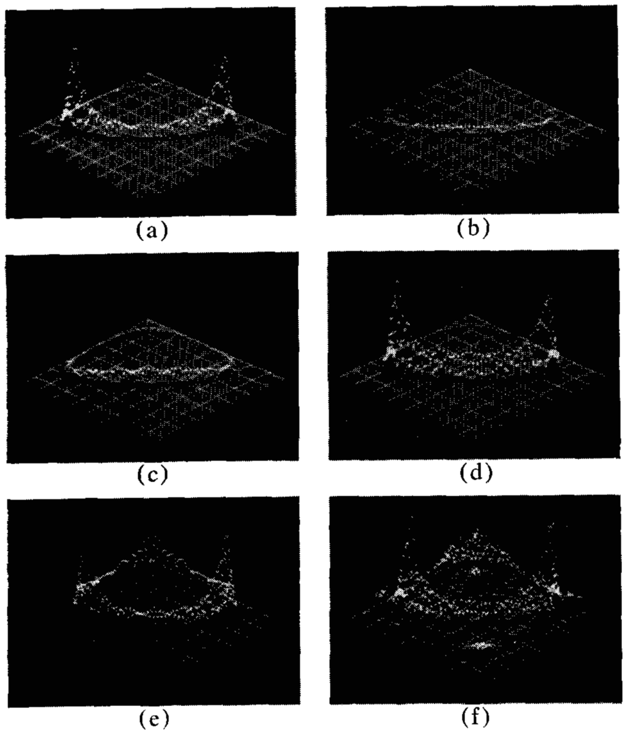

Figure 2 Use of displays to fit two-dimensional data. The individual terms and factors are adjusted to give the best-appearing fit to the data.
(a), (b), (c): Three independent terms $T_n(i, j)$ in the theory with coefficients C_n adjusted to fit the data.
(d): Sum of these terms for each i and j: $C_1T_1(i, j) + C_2T_2(i, j) + C_3T_3(i, j)$.
(e): Final theoretical result with over-all coefficient, of the form $C_0[C_1T_1(i, j) + C_2T_2(i, j) + C_3T_3(i, j)]$.
(f): Experimental results for comparison.

Figure 3 Integration of a peak using the light pen. In (a) the peak of interest is indicated by tagging channels on either side with the light pen. The system responds by blanking those channels. The quadratic fit to the blanked-out channels is displayed in (b). If channels part way up the peak are designated as background (c), the erroneous fit is obvious (d) and the operation can be restarted.

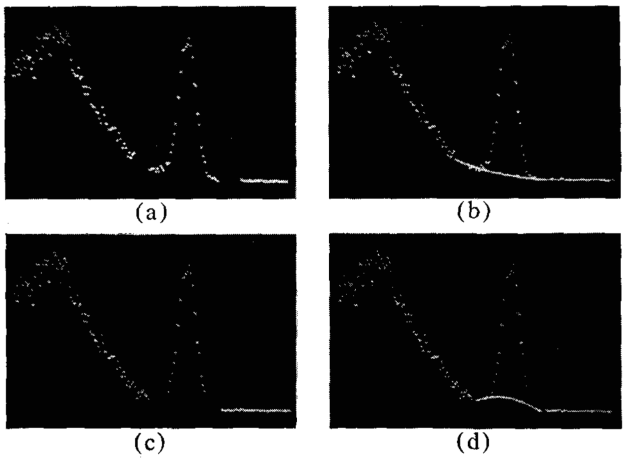

Interactive data analysis

The multidimensional experiments have profited enormously from the existence of on-line computer facilities—indeed we would not have attempted most of them otherwise. Nevertheless they represent a minority of the experiments in progress on the accelerator at any one time. The advantages provided by the computer system to the more common one-dimensional experiments have been in convenience of use, efficiency of operation, and speed of data analysis.

In one-dimensional spectra like that in Fig. 3, one is usually interested in the area and position of one or more peaks. These are measured using a light pen to delineate the peak on the displayed spectrum as illustrated. The background on either side of the peak is fit to a second-order curve and interpolated under the peak. Careless use of the light pen may give erroneous fits, but since the computer immediately displays the results obtained, the experimenter can detect errors and repeat the operation. The assumed background under the peak was first displayed during an experiment which had a difficult background to fit; the rejection of obviously unrealistic fits resulted in a factor of two reduction in the scatter of the area measurements.[3]

The close partnership of the computer and experimenter in our laboratory generally has been found preferable to a more distant relationship with a powerful computer-center system. If the object is to increase the productivity of the laboratory personnel in their research, then ease of access is more important than computing speed or memory size. Figure 4 shows the fitting of theoretical parameters to data by R. G. Van Bree and E. J. Schneid.[4] With the aid of the display and the small computer they were able to analyze their data in much less of their own time than if they had used a large remote computer. This kind of operation is also attractive from a cost standpoint since the hourly rental for the small computer is quite low.

Most data acquisition and analysis programs for our computers have been written in FORTRAN. Only the I/O routines for devices unknown to FORTRAN have been hand-coded: the nuclear data interface, the display, the light pen, and the plotter. Of the 8K words of core, data arrays usually occupy 2K, hand-coded routines perhaps 0.25K, and FORTRAN object programs and run-time systems the remainder. For two-dimensional experiments, FORTRAN must be dispensed with; up to 6K are used for data.

Saturation: the penalty of success

Within six months after its installation at the nuclear physics laboratory, all experimental groups were using the computer for data collection. Since the accelerator runs 24 hours a day, no time was left for data analysis and program development. This bottleneck was solved by the addition of a program-compatible SDS 925, with a line printer, card reader, and plotter that is switchable between

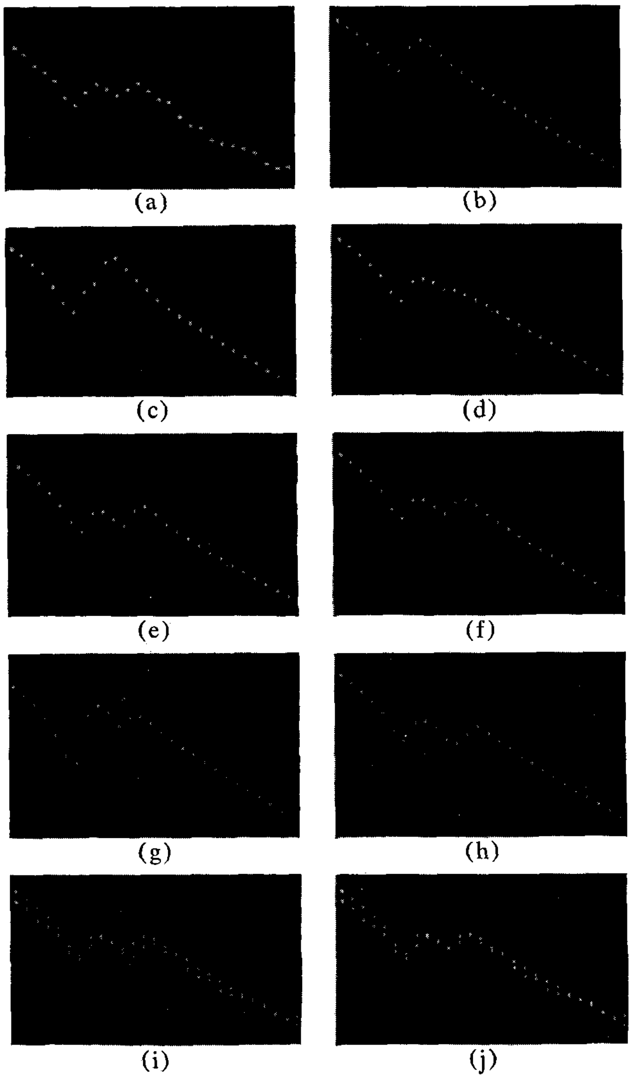

Figure 4 Interactive fitting of five resonances to data from the reaction. The approximate positions of the resonances are initially known but their exact positions, strengths and spins must be obtained from the fitting process.

(a): 50 channels of data comprising the region of interest.

(b)-(f): Five resonances are added, one at a time. The parameters used here were those which gave the best final fit.

(g): Similar to (f) but the spin of the first level is assumed to be 3/2 rather than 1/2.

(h): Similar to (f) but the fourth resonance is assumed 50 keV too high.

(i): Spectrum (h) and data (a) superimposed for detailed check.

(j): Spectrum (f), the final fit, and data (a) superimposed.

it and the 910, as in Fig. 1.

The dual computer system has worked well for two years, but the increased power of the system has been overtaken by the demand. The 910 is now used exclusively for data collection, while the 925 does data analysis, program development, and some theoretical computation. However, the 925 can serve only one user at a time, and frequently this user is only contemplating the displayed spectrum, light pen in hand. While this procedure is more productive than earlier methods of data analysis, it is obviously far from optimum. At the very least the computer should be made productive while the light pen user is studying and tagging his spectra.

The new generation

Rather than continue to add to the present system, it was decided to build an entirely new system. The new system is an attempt to provide the best possible assistance in the conduct and analysis of experiments. In this effort we have been guided by our experience with the current system in choosing which operations to emphasize with special hardware. On the basis of past experience, it is likely that data collection rates, the number of channels* desired, and the dimensionality of experiments will increase. The more the limits on these factors are raised, the longer the system will be satisfactory to its users.

Goals of the new system

In designing the new information handling system, several goals have been set:

1. Simultaneous usage by at least two data taking groups and two or more analysis groups;
2. higher I/O reliability;
3. greater ease of use; and
4. greater computing power.

Simultaneous usage of the system by a reasonable number of people (relative to the size of our laboratory) has become possible because of reductions in the costs of logic and of rotating storage. The displays can be regenerated from a disc rather than from computer memory, freeing the computer for other operations. Users can thus reflect at leisure on the significance of their displayed data without holding up other activities. Unfortunately, the computer must spend much of its newly available free time in the servicing of data stored on the disc. Since the disc does not permit random access as in core storage, the updating of a few events requires more complex processing.

The solution to this problem and the key to high productivity in the new system is the increase in the number of linked processors that operate simultaneously. We can use a small computer for data acquisition and disc updating, under the control of a main computer which is largely free for analysis and theoretical work.

* "Channel" is used here in the sense of bins into which data are sorted rather than to mean an I/O path into a computer.

In addition, I/O processors will control the flow of data in and out of the memory of both computers, permitting I/O operations to be overlapped with computation. The data/display disc is attached to a second memory bank on the data acquisition computer; once a transfer is started, it will proceed at a rate of 10^6 bytes/second while the computer returns to data sorting, and new data come in through the I/O processor.

Data from the experiments will come in through two control and formatting units designed at Brookhaven National Laboratory by P. C. Rogers and G. E. Schwender.[5] This interface will permit easy attachment of numerous analog-to-digital converters and other devices, resolve priorities, and take full advantage of the I/O processor system. The availability of two interfaces will permit one experiment to be set up and checked out while another is in progress.

In controlling this complex set of interconnected processors, we differ from most large/small laboratory systems. The small machine (Sigma 2) will operate as a satellite under the control of the larger computer (Sigma 5) and as a programmable front end for data collection and disc updating. Thus the user, whether collecting data, analyzing it or debugging programs, will have access to all peripherals and to the full monitor system of the large computer. As indicated in Fig. 5, programs for the Sigma 2 will be called from the Sigma 5 disc and sent to the Sigma 2 by typed commands. Data will be sent back for storage, analysis, and output.

This type of operation requires some degree of time sharing. For the small number of users on the system (4 or 5) we can achieve an adequate level of time sharing simply by swapping the main computer core between users. Data acquisition, analysis and nonconsole background jobs can be swapped equally easily because the critical real-time tasks are not swapped but reside in the Sigma 2. The displays, once formatted, are regenerated from a separate disc and require no CPU time.

With the time sharing facility we expect to lower the user's programming barrier substantially. Nuisance operations with cards and paper tapes will be eliminated by disc storage of source language programs coupled with editing and debugging facilities. The card reader and line printer fit naturally into the system and provide greater I/O flexibility than is possible on systems where time-sharing users have access only to teletypes. Two magnetic tape drives will also be available for storage beyond the capacity of the disc.

Running the satellite computer under the control of the main computer provides a number of advantages not otherwise realized. The maintenance and documentation of data acquisition routines will be centralized since these routines will be kept on the main computer disc. Likewise, loading and running of these routines will require only a load command to the monitor. Once a data acquisition task is running in the smaller computer, the experimenters can analyze data from a previous run at the same console, checking the progress of the data taking at will. No I/O gear will be needed for the satellite computer save the intercomputer connection and the display disc, reducing the probability of downtime.

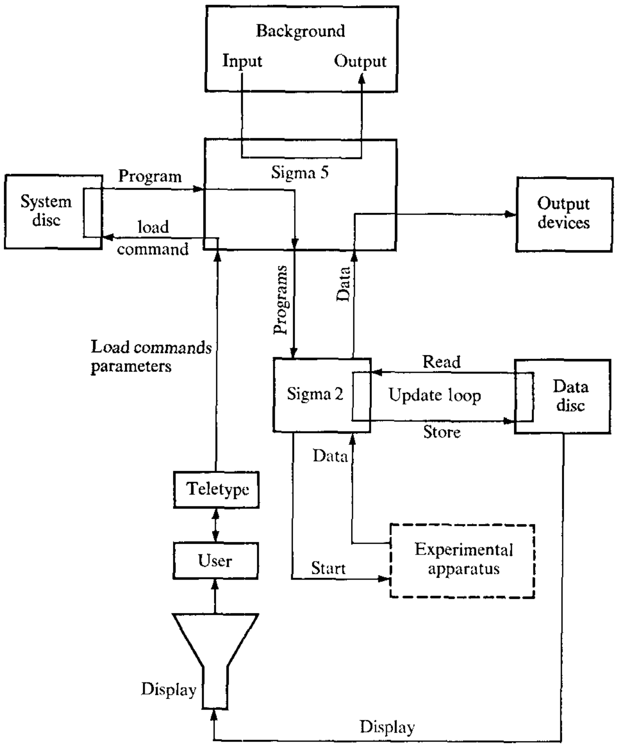

Figure 5 Flow of control and information. The user types his commands to the Sigma 5 monitor, which loads programs from the system disc. Analysis programs are executed in the Sigma 5 with display information sent over to the Sigma 2. Data acquisition programs are sent to the satellite for execution; data are returned to the main computer for analysis and output.

The use of time sharing will permit convenient program editing and debugging facilities. The time spent in the debugging stages of program development will be cut substantially, hopefully by an order of magnitude.

Potential problems

In introducing the new computer system, we will be asking the users to cope with a system of much greater complexity than the previous one. They will be able to do more with it, but at the cost of learning new procedures. The computer must be set up as a perfect servant, unobtrusive yet assignable to any task with a minimum of instructions. The physicist should be able to concentrate on problems of physics. If he has the time and the inclination he can write intricate and specialized machine language programs, but he must be able to perform experiments without becoming a computer specialist.

In order to accomplish this, a systems programming group will be necessary for the first time. The responsibility for writing the machine-language data acquisition and disc-updating programs will fall on this group. Our experience with the original system indicates that a very small number of such programs will suffice for all users. Having these programs written, checked out fully, and stored in the library will minimize interuser interference and reduce the user's programming load.

Compatibility with the present system in the transition period will require much attention. Although the present system is overloaded, it works well for the user once he has access to it. A new system with good accessibility but no programs available would be no improvement. For short-term use, compatibility will be maintained by two features: The FORTRAN main programs will be handled by the new compiler, and the machine language subroutines replaced by identically called routines written by the systems group. Experimenters will be able to expand their usage to take advantage of the new features at their own convenience.

Light pen operations will be more complex with the new system. For our present core-generated displays, the light pen interrupt can set a tag bit on each data point indicated by the user. This cannot be done when one displays from a disc. If the light pen information is used like experimental data to update the disc, a time lapse occurs before the display responds by blanking tagged points. A simulation study indicated that for a two-dimensional display the light pen information should update the display at least four times per second. This is hard to guarantee if data are coming into the system at high rates; other methods of handling the light pen are being considered, such as a small core memory for the light pen tag bits only.

Conclusions

The new nuclear physics laboratory computer is well out of the class of small computers. An installation that started with an 8-μsec cycle computer, with 4096 24-bit words of core, is growing in complexity to two linked processors and a total of 152K 8-bit bytes of 0.9 μsec memory, and two discs. There is no question that the computer power will be adequate for some time to come; the difficult tasks will be in the management of these resources to make them maximally available to the users. If we are successful we will have an excellent framework for innovations in the conduct of computer-aided experiments. While the availability of two computers may not be duplicated in every laboratory, the experience with our large-small computer combination may be applicable in environments where one or more central computers operate in close partnership with satellite data-taking computers spread over a campus or research institute.

Acknowledgments

Portions of both the original and new computer systems described here have been supported by grants from the National Science Foundation to Rutgers University.

References

1. J. V. Kane, *Proc. Conf. on the Utilization of Multiparameter Analyzers,* 149, Report CU(PNPL)-227, Columbia University, New York (1962).
2. D. P. Boyd, Ph.D. Thesis, Rutgers University (1968).
3. P. F. Donovan, J. V. Kane, J. F. Mollenauer, D. P. Boyd, P. D. Parker, and Č. Zupančič, *Phys. Rev.* **158,** 973 (1967).
4. M. Wiesen, E. J. Schneid, and G. M. Temmer, *Bull. Am. Phys. Soc. II,* **13,** 563 (1968).
5. G. E. Schwender and P. C. Rogers, submitted to 15th Nuclear Science Symposium, Montreal, October 1968.

Received August 7, 1968