OAK RIDGE NATIONAL LABORATORY

OPERATED BY MARTIN MARIETTA ENERGY SYSTEMS, INC.
POST OFFICE BOX 2008, OAK RIDGE, TENNESSEE 37831-6285

ORNL
FOREIGN TRIP REPORT

ORNL/FTR-3405

DATE: October 9, 1989

SUBJECT: Report of Foreign Travel of Robert L. Varner, Jr., Physicist, Physics Division, September 3-30, 1989

TO: Alvin W. Trivelpiece

FROM: Robert L. Varner, Jr.

ORNL/FTR--3405

DE90 001545

Purpose

(a) To confer with staff members at CERN, Geneva, Switzerland, about data analysis and acquisition hardware and software.

(b) To attend the Nuclear Structure in the Era of New Spectroscopy Workshop (Section A - Data Processing and Correlation Analysis), Copenhagen, Denmark.

Sites Visited

Sept. 5-6, 1989	UA1 at CERN DD Division at CERN Geneva, Switzerland	J.-P. Porte R. Brun
Sept. 11-29, 1989	Workshop Copenhagen, Denmark	A. Holm

Abstract

This trip had two purposes. The first was to attend a workshop on data acquisition and analysis of data from highly segmented Ge gamma-ray detectors. The systems discussed were mainly for a detector project in Europe (Euroball), but the principles and many of the details are the same as for the proposed GAMMASPHERE detector to be built at ORNL. I also visited CERN to discuss the MICRON/MACVEE system for controlling CAMAC and VME-based electronics from a Macintosh personal computer. I discussed the system with one of the implementors and received a copy of the latest version of the development software. This software will be used at ORNL to develop small, dedicated data acquisition systems.

Introduction

My trip had two major components, a visit to CERN to discuss software and hardware for data acquisition and analysis, and participation in a workshop on data acquisition for highly segmented, high resolution gamma-ray detectors at the Niels Bohr Institute in Copenhagen, Denmark. I will discuss these two parts separately. The workshop was organized into two parts, invited talks on topics related to the workshop and work of study groups on the problems of the acquisition and analysis. I will discuss these separately.

Discussions at CERN

Members of the Physics Division at ORNL are looking for means to assemble flexible, inexpensive and small data acquisition systems for testing and monitoring of detector systems such as the BaF_2 array. One clearly useful choice for this is the Apple Macintosh computer coupled with the MACVEE/MICRON interface to VME and CAMAC. This board was developed at CERN by the UA1 collaboration to control and monitor the data acquisition and playback system for the entire UA1 experiment. The hardware is commercially manufactured, but there is no commercial source of the software. The head of the software group for UA1 suggested that the best way to obtain the software was to come to CERN and talk with members of the development group. That is what I did. I met with Jean-Paul Porte and Sergio Cittolin for several hours, during which they demonstrated the functionality of the Macintosh-based control and monitoring system. The system can monitor data processing within VME systems and control the data source for the data processing system, which can be the on-line experiment, high-speed IBM data cartridges, or EXABYTE tape. UA1 uses the system entirely as the user interface to software running in VME modules (Motorola 680x0 processors or CERN 3081 emulators) and a display of data or status information from the ongoing acquisition. The software is also used to compile and link code for downloading the 680x0 VME processors.

After these demonstrations, I was shown the software system, called MACUA1, which includes a command processor, a high-quality Fortran compiler, a code management system and numerous support libraries. These libraries include basic VME and CAMAC control routines, large portions of the CERN KERNLIB and a graphics library. This software is all in the public domain, and is shared by UA1 with members of the physics research community. I obtained a copy, which fits on one Macintosh diskette, as well as a printed copy of the documentation. In addition to the software I have already described, I received a copy of a newly developed CAMAC list processing utility, Macintosh communications software and EXABYTE tape drive interface.

The UA1 group maintains a list of users of the MACUA1 system and will soon make updates to the software available on the world-wide HEPNET from a VAX at CERN.

The detailed description of the software and its function given to me by the members of the UA1 group, as well as the demonstration of the flexibility and power of the Macintosh applications, will benefit the Physics Division greatly in the early stages of development of an acquisition and control system for the BaF_2 project and others that will follow later.

I took advantage of this CERN visit to discuss with Rene Brun of the CERN DD division the future direction of PAW (Physics Analysis Workstation), a graphical data examination and processing software package which runs on VAX, Apollo and Hewlett-Packard workstations. In particular, I was interested in whether or not PAW would work with X-Windows, a new standard in workstation interfaces. We discussed several related issues as well.

The PAW group is currently working on X-Windows interfaces for PAW, but that work is still in its early stages. There are problems with the X11 standard which severely reduce the efficiency of PAW when displaying most physics data, especially when the medium of communications is the network. X11 has no concept of histograms in the server software, which means that PAW must assemble histograms in terms of the line drawing facilities provided. This can increase the data load of a local area network by a factor of 4 to 5 and seriously degrade the response time of PAW to user commands. Brun is most adamant about quick response time, so that the physicist using the product does not feel he is waiting. CERN is currently negotiating with the X11 standards group to add these features. In the meantime, they are experimenting with various other means of distributing the work of PAW, in particular through the use of data servers, rather than graphics servers. Such a distribution system requires more local intelligence than, for example, an X-windows terminal, but only as much as a typical personal computer which is a network client. Such efforts are much farther along, and are all based on using TCP/IP as the communications protocol. The software, as it exists, will be available to us over the HEPNET, as well as some assistance with porting it to Physics Division computers.

As part of our discussion, Brun mentioned standards and likely directions for future computing at CERN. CERN is very aware of the coming wave of standards, POSIX, OSI and X11, and the growing power and decreasing price of RISC-designed processors running UNIX, and they are glad to embrace these standards. Brun is of the opinion that VAX/VMS will be "dead in two years" as an operating system of choice for scientific applications. We should watch carefully the direction CERN and other major labs in nuclear and particle physics go, to be best able to take advantage of the software and hardware they produce.

Workshop at Copenhagen

The workshop at the Niels Bohr Institute in Copenhagen was titled "Nuclear Structure in the Era of New Spectroscopy, Section A - Data Processing and Correlation Analysis." It was intended to discuss the physics and engineering possibilities of highly segmented, 4π solid angle Ge detector systems. It was organized by Anders Holm of NBI, who intended it to be more slowly paced than is customary in workshops, to give the participants time to interact and achieve a "common vocabulary" for data acquisition systems and a common understanding of the function and goals for such systems. The meeting consisted of invited talks and other presentations in the mornings, with study group discussions following lunch in the afternoons. I will summarize the invited talks, then discuss the results of the study groups.

There were three kinds of talks presented: (1) talks about existing data acquisition systems, (2) talks about the physics to be gained using these new kinds of detectors, and (3) lectures on computer science and software design.

The first speaker was D. Hensley of ORNL, who discussed the components of acquisition, as well as the problems and nonproblems of designing acquisition for a Euroball or GAMMASPHERE detector. In two talks about data acquisition systems, M. Maier described the acquisition system for the 4π detector at MSU, a system based on CAMAC FERA ADC's processed by VME computers, and the acquisition system described in the GAMMASPHERE proposal. J. Hansen of NBI described the data acquisition system for the ALEPH detector at the CERN LEP accelerator. The project began in 1983 and designed several hardware and software components for use in the acquisition system. The talk emphasized the need for detailed planning before beginning code development (they spent 1.5 years simply planning), the use of higher-level languages (such as C), and the need for high quality, accessible documentation. It also emphasized the need to use as much available software as possible, and to stick with early decisions. In the last week of the workshop, R. Lieder of Julich and F. Beck of Strasbourg described the current status of the Euroball project and the data acquisition design and development in particular. The Euroball collaboration has set an ambitious goal of having the detector operating in Phase I by January 1992. They also have the more ambitious goal of having the signal processing electronics integrated onto VXI cards in about 18 months. M. Maier was skeptical that the integration goal could be achieved in the desired time. This issue was discussed in more detail in the study groups.

In addition to the above talks, talks about local data acquisition systems were given by M. DePoli (Legnaro), J. Poggioli (Saclay), C. Ender (Heidelberg) and G. Vedovato (Legnaro). These systems were based on various front-end busses and used distributed processing to accomplish the readout and data filtering functions. The Heidelberg system was notable in that it uses a 680x0-based system designed and built entirely at Heidelberg. The computer system was the POLYP system described earlier by Ender, which performed several levels of filtering in the event stream to reduce it by 2 to 3 orders of magnitude for tape storage.

The physics talks presented generally emphasized the variety of new data available from these new detectors, and the difficulty of analyzing that data. The biggest gain is, of course, the much higher efficiency for many-fold gamma ray coincidences with high resolution. D. Radford of Chalk River discussed methods he uses to analyze high spin data taken with the 8π spectrometer at his laboratory, especially 3-fold gamma data. His analysis technique emphasizes using lower-fold data to enhance his understanding of higher-fold data, for example, using 1-fold data to identify peak energies. He also uses model information from physicists to analyze 2- and 3-fold data, by looking for coincident weak transitions with predicted energy separations. This technique has allowed him to rapidly identify superdeformed bands in some nuclei. In his talks he emphasized the need for using the computer to save the physicist tedious work, such as recording details of level schemes, and to allow physicists to quickly test likely level schemes against data. Some of this has been implemented in a program at Chalk River called "LF8R." He emphasized also that he believes that little is to be gained from analyzing greater than 3-fold data directly in histograms. Rather, the data may be looked at in 3-dimensional projections, then used in event lists to test higher-fold coincidences. The major problem of higher-fold data is the storage requirement of n-dimensional data, and the representation for the physicist.

B. Herskind of NBI talked at length of the need for complete spectroscopy of the nucleus, and of the difficulties involved in identifying superdeformed and hyperdeformed bands in candidate nuclei. The much higher resolving power of this new generation of detector encourages him to believe that these bands can be found. T. Lonnroth of Finland reported on work he is doing, using a 100-element Transputer array to process 3-fold gamma data. The computer gives him a performance of 150 MFlops and permits him to perform unrestricted analyses of the 3-fold data, in the sense that he does not start from lower-fold data to interpret the higher-fold. W. Urban of Bonn discussed his techniques for representing data (as compressed event lists) and for essentially statistical background subtraction performed during the event sorting. The technique generated much discussion about its validity. A. Lampinen of Finland discussed work he has done in representing higher-fold matrices, using data structures borrowed from robotics called OCTREES, which allow an easily searched and correlated representation of three-dimensional spaces, as in 3-fold matrices. The advantage of this is that octrees are well understood, and much software exists to manipulate them. It is not clear that the representation is compact enough to be useful for higher-fold data.

In general the computer science talks seemed intended to present a means of thinking about programming different from the average physicist's way of thinking. These talks were begun by a physicist-turned-computer scientist, P. Moller-Nielsen of Aarhus University, who discussed the state of computer science (its art, craft and logic) and the folklore of the field. He emphasized in particular the transition of computer science from pure art and folklore to software engineering and the current state of testing and verification. Later he talked about the production of ASIC's, Application Specific Integrated Circuits. The necessary tools exist to permit the design of ASIC's even in a small laboratory, and the production process is a fairly standard service provided by the semiconductor industry.

In another series of talks, U. Caprani, also of Aarhus, discussed techniques for modeling the software of multiprocessing systems. He discussed synchronization, modularity of design, communications, and even the application of these ideas in a real time environment. These talks were quite valuable to me, providing a very useful paradigm for thinking about multiprocessing systems.

Other speakers included K. Conradsen of NBI, who spoke on image processing techniques, in particular the extraction of quality images from relatively noisy data; C. Ender, who described the POLYP multicomputer, a 100 MIPS multiprocessor array at Heidelberg, constructed and designed entirely at the University; and B. Lautrup of NBI, who spoke on the general properties of neural networks and software models of them. I spoke briefly on the need for standardization in the development of software, to enable data acquisition and analysis systems to keep up with the rapid pace of hardware technological advances.

Workshop Study Groups

There were two study groups in the workshop, one concerned with data acquisition; the other with data analysis. I participated in the discussions of the data acquisition group, led by D. Hensley and M. Maier, which considered the data acquisition system from the front-end electronics to the final event stream to be stored by the user/experimenter. Many details were considered, including the complexity and flexibility of the system, the scale of integration of the electronics systems, and even the desirable functions of the components, including computer control. Much of the discussion of the workgroup involved understanding desires of the potential users for not only high resolution gamma rays, but also for high quality calorimetry on the part of the ball, which significantly increased demands on the processing electronics in each of the Ge detector processing channels. The result of the study group was a design, in principle, of a data acquisition system for any highly segmented, high resolution gamma detector. Each major system of the design has potential for significant integration in the channel and trigger electronics, should manpower and money be available for the research, as well as for significant parallelism in the readout control and event formatting. This design can be used very much as is to guide the GAMMASPHERE effort at ORNL in the construction of the data acquisition system.

The data analysis group presented reports at several times during the workshop (as did the acquisition group). In their final summary, they discussed many points related to analysis software development and made several recommendations. Foremost, there is a need for standardization of some of the technical vocabulary in the processing of high-fold gamma ray data. Several conference members were using the same words to refer to very different aspects of data processing. Next, there is a need for standardization in nuclear physics, to make the variety of software being developed transportable. Some of the standardizations included histogram formats, event formats, user interface and graphics standards and adoption of industry standards. Some suggested that certain vendor hardware and software could be chosen, but most workshop members (the vocal ones, at least) felt that vendor-independent standards were the only serious ones to accept.

Much of the rest of the analysis summary dealt with the enormous problem of storage and examination, that is, how to look at high-fold data in reasonable time without using hundreds of Gigabytes of storage. Some of the ideas include storing data as simple events, which occupy much less space than n-fold histograms would, where n is greater than 3. There were ideas about histogram storage, using dense storage techniques and data compression without sacrificing information by using nonlinear energy transformations for histograms, which could reduce 3-fold histogram sizes by a factor of 10.

More recommendations by the analysis group concerned the amount of data and the physicist's ability to understand it. The group suggested that such analysis requires a high level of interaction between the physicist and the software, with the ability to propose level schemes, test them and discard failed schemes, much as one interacts with personal computer spreadsheets. Some members of the analysis group believe that analyzing fold greater than 3 data is not possible, at least not in the form of n-dimensional matrices. This is a question to be resolved after the first higher-fold data become available, perhaps.

Conclusions

This meeting proved invaluable to me. I learned much about computation and data acquisition, in an atmosphere where it was possible to think about the presentations, instead of just absorbing data, as might have occurred at a conference. I believe that the discussions there can provide valuable guidance to GAMMASPHERE, especially since several potential collaborators in the GAMMASPHERE development were participants. In particular, I think that contacts we made with representatives of the Euroball signal processing and monitoring groups will provide valuable collaborations for GAMMASPHERE.

The only disquiet I feel about the workshop is (1) that so few members of the Euroball collaboration came, except to give short summary talks, and

(2) that so few potential users of the device came to express their needs. In the absence of the users, it would have made more sense to schedule this section of the workshop after the more physics-oriented sections (B and C) had concluded, to provide better direction for our discussions. As far as the lack of Euroball attendance goes, one of the major reasons for this workshop (according to Anders Holm) was to offer the Euroball collaborators a chance to resolve problems of communication and to think about the system, rather than technical solutions to small problems. It may be that we shall see difficulties in the timely execution of the Euroball project.

Appendix

A. Itinerary

1989

Sept. 3-4	Travel from Oak Ridge, TN, to Geneva, Switzerland, via plane
Sept. 5-6	CERN, Geneva, Switzerland
Sept. 7-9	Vacation and weekend
Sept. 10	Travel from Geneva to Copenhagen, Denmark, via plane
Sept. 11-29	Workshop at Niels Bohr Institute, Copenhagen
Sept. 30	Travel from Copenhagen to Oak Ridge, via plane

B. Persons Contacted

J.-P. Porte	CERN, Geneva, Switzerland
R. Brun	"
S. Cittolin	"
A. Holm	Niels Bohr Institute, Copenhagen, Denmark

C. Computer Software Acquired

MACUA1 Software System (1 Macintosh diskette and printed copy)

CAMAC List Processing Utility, Macintosh Communications Software and EXABYTE Tape Drive Interface