OPERATIONS

Cyclotron Operating Experience

D.R. Poe and H.A. Thulin

The first beam was extracted from the K1200 cyclotron on June 6, 1988. Initial operation indicated that the K1200 is a well understood accelerator; beam is obtained with the computed parameters and only minor changes are necessary to optimize the extracted beam. Operation is reproducible, with beam obtained at identical settings after intervals of several weeks.

In October, the beam was transported through a superconducting beamline to the interim experimental area, which contains the general purpose 92" scattering chamber and the 4π array. Although this interim area is very near the cyclotron, and there is little magnetic analysis, the beam has sufficiently high optical quality for demanding experiments to be done.

TABLE I: K1200 experiments in 1988.

Month	Expmt #	Spokes-person	Beams Used	
October	Disc.	Stevenson	$^{14}\mathrm{N}$	75 Mev/u
December	86022	Gong	$^{4}\mathrm{He}$	20 Mev/u
	"	"	$^{4}\mathrm{He}$	30 Mev/u
	86040	Gong	$^{129}\mathrm{Xe}$	31 Mev/u

TABLE II: K500 time distribution.

Use category	Hours	Fraction of operating period
Operation		
Research	1803.00	50.0 %
Development	166.50	4.6 %
Overhead	1048.25	29.1 %
	3017.75	83.7 %
Maintenance	47.25	1.4 %
Breakdown	538.75	14.9 %
TOTAL	3603.75	100 %
Off (shutdown periods)	5180.25	

TABLE III: K500 extracted beams.

Ion	E/A [MeV/u]	Hours	% Time
$^{2}\mathrm{H}^{1+}$	53.0	218.5	9.4
$^{4}\mathrm{He}^{1+}$	13.0	23.5	1.0
	25.0	17.5	0.8
$^{4}\mathrm{He}^{2+}$	54.0	7.5	0.3
	60.0	1.0	0.0
$^{12}\mathrm{C}^{3+}$	15.0	21.5	0.9
$^{12}\mathrm{C}^{4+}$	30.0	20.3	0.9
$^{12}\mathrm{C}^{5+}$	50.0	278.8	12.0
$^{14}\mathrm{N}^{2+}$	7.2	3.0	0.1
	8.0	1.3	0.1
$^{14}\mathrm{N}^{3+}$	17.0	29.5	1.3
	20.0	1.0	0.0
$^{14}\mathrm{N}^{4+}$	22.4	9.8	0.4
$^{14}\mathrm{N}^{5+}$	30.0	25.3	1.1
	35.0	112.3	4.8
$^{14}\mathrm{N}^{6+}$	45.0	45.8	2.0
	50.0	21.5	0.9
$^{16}\mathrm{O}^{6+}$	35.0	98.2	4.2
$^{16}\mathrm{O}^{7+}$	50.0	35.3	1.5
$^{20}\mathrm{Ne}^{7+}$	35.0	31.5	1.4
$^{20}\mathrm{Ne}^{8+}$	45.0	89.5	3.9
$^{22}\mathrm{Ne}^{3+}$	5.1	119.5	5.1
$^{36}\mathrm{Ar}^{10+}$	20.0	16.5	0.7
$^{36}\mathrm{Ar}^{11+}$	35.0	373.8	16.1
$^{40}\mathrm{Ar}^{4+}$	4.0	48.5	2.1
$^{40}\mathrm{Ar}^{5+}$	4.8	143.8	6.2
$^{40}\mathrm{Ar}^{7+}$	9.0	32.5	1.4
	10.0	21.0	0.9
$^{40}\mathrm{Ar}^{9+}$	20.0	38.5	1.7
$^{40}\mathrm{Ar}^{12+}$	35.0	165.0	7.1
$^{56}\mathrm{Fe}^{8+}$	7.2	156.3	6.7
	8.0	4.0	0.2
$^{86}\mathrm{Kr}^{10+}$	6.0	1.5	0.1
$^{86}\mathrm{Kr}^{15+}$	10.0	32.0	1.4
$^{86}\mathrm{Kr}^{17+}$	15.0	46.3	3.3
		2321.0	100.0

The K1200 was commissioned using the compact ECR source, and the first experiment, an observation of high-energy gamma rays produced by 75 Mev/nucleon, $^{14}\mathrm{N}$ ions, was performed with this source. Following the initial experiment, the room-temperature ECR source was switched to the K1200 and experiments with more difficult beams began. Table I lists the beams used for the K1200 experiments carried out in 1988.

During the next year, development efforts will be aimed at increasing the energy and reliability of available beams and improving the injection efficiency which, because of the influence of fringe fields, is now low when the K1200 is operated at high magnetic fields.

Operation of the K500 continued during 1988, although construction conflicts greatly limited K500 use. The total number of hours of operation was 3018, compared to 4368 in 1987. The reliability of operation, defined as the percentage of total scheduled operating time for which the operations were continuing, or (Research + Development + Overhead)/(Total - Breakdown - Maintenance) was 84.9%. This is the highest value of reliability obtained to date. Details of the K500 time distribution are given in Tables II and III.

Supplying Cryogens -- 1988 Update

H. Laumer, R. Fontus II, N. Friedman, J.A. Nolen, D.E. Pendell, R. Selden, and R. Zarobinski

Cryogens to operate the K1200 and K500 magnets, and the cryopanels in the beam chambers of the associated cyclotrons have been supplied by a cryogenic fluids distribution system described in the 1987 Annual Report. During this reporting period, two extensions to this system were implemented. One branch out of Box #2 supplies the four beamline magnets needed to conduct experiments in the 92-inch chamber and the 4π array in their present positions. A second branch near Box #1 leads to a new control box supplying cryogens to the 6000-l-capacity test dewar. Finally, the line which will feed most of the beamline magnets has been extended through the old K50 vault wall into the future transfer hall.

The operating mode of the beamline magnets is a departure from the continuous direct feed for cyclotron magnets and cryopanels. The beamline magnets are batch-fill magnets. The various magnet reservoir dewars allow them to operate from one to over six days between refills (details elsewhere in this report). At this time they are usually all refilled once a day because the prototype ± 16° dipole requires this frequency. It takes from 150 l to 190 l of liquid helium per day to refill two doublet and one singlet quadrupoles as well as one dipole.

Initially, we encountered some difficulty in that the filling procedure indirectly affected the liquid helium supply to cyclotron coils and cryopanels. Undesirable side effects were temporary lowering of coil cryostat liquid helium levels and warming of cryopanels, which in turn caused vacuum deterioration. An automatic filling cycle was developed. This involved converting a few manual valves to pneumatic operation on Box #2.

The supply valves on the beamline magnet cryostats are electrically operated. Filling can be manually supervised from a switching console. In the automatic filling mode, relay logic makes decisions based on the state of one timer, two temperature sensors and two level sensors. Upon a signal from a Modicon process controller, a liquid helium supply valve opens partially to cool down transfer lines for a timed interval. After this interval, the liquid supply valves to individual dewars open, provided the level sensors indicate levels below their makeup setpoints. As the return gas temperature drops below a preset value, the gas is routed to the cold gas return stream to conserve liquid. At this time, the liquid supply valve is opened more to compensate for the reduced driving pressure.

It is the careful balance of this sequence that minimizes adverse effects on the other devices. Only one quadrupole level sensor, manually preselected, is monitored during the filling cycle. When its level rises above the high set point, the liquid supply to all three quadrupole cryostats is cut.

In practice, the quadrupole with the fastest boil-off rate is selected to control the filling cycle. The dipole, which is refilled most frequently, has its own level sensor control. It has the longest filling time and supersedes the other level sensor, in that the main liquid helium supply valve closes when the dipole is full. At this time, the boil-off gas is switched to the 80 K gas-return circuit. By this means the cryostat pressures are stabilized and lead flow is stable in turn. The system is then ready to restart a filling cycle initiated by a manual or Modicon signal.

The liquid nitrogen dewars in the beamline magnets have a longer holding period between refills; however, in practice they are filled daily. There are no level sensors in these dewars. The dewars are filled in series with one supply valve and one exhaust line. Boil-off gas has to follow this same route. The dewars are well-baffled, so that no liquid appears in the exhaust until the last vessel is full. A simple device detects when liquid appears in the exhaust line. A 30-cm-long capped tube, installed with a "T" in the exhaust line, acts as a phase separator. A thermocouple is mounted on this cap. The temperature remains near room temperature until liquid nitrogen drops into the capped portion of the "T" and lowers the temperature suddenly. A simple nonindicating thermocouple controller then causes the supply valve to close, waiting for the manual or Modicon signal to start the next filling cycle.

The operation of the cryopanels has been simplified for the cyclotron operators. By modifying manual gas return valves to pneumatic action, cool-down is governed by Modicon-controlled switching circuits. Operators select a page on the Modicon's Panel Mate touch screen and push "Start" for liquid helium and liquid nitrogen cooling to commence as desired at K500 or K1200 cryopanels. Temperature sensors on the warm-gas-return lines determine when to switch to cold helium gas return. Upon shutdown by touch-screen command, a timed sequence under Modicon control avoids trapping liquid and finally sets valves to be ready for the next cool-down cycle.

The CTI-1400 liquefier supplied liquid helium to the K500 coil and cryopanels during the initial years of K500 operation. It has now been moved from the east high bay to the west high bay to near the S800 spectrograph, with the intention of eventually operating the spectrograph magnet with this refrigerator. A liquid nitrogen supply line, as well as high-pressure gas and compressor suction return lines, had to be extended 300 ft to the new location. The temporary transfer lines to the K1200, which had been replaced, were recycled in this construction. The CTI-1400 is presently operated as a magnet test facility for beamline magnets. It also has been supplying liquid to the medical cyclotron.

The 6000-l test dewar has also been used to perform coil tests in this operating period. Its operation is usually routine, but its big liquid consumption load sometimes meant it could not be operated at will, but had to be scheduled to allow liquid inventory to build up.

The record for keeping devices supplied with cryogens was good. There were a few compressor and refrigerator equipment failures that resulted in lost operating time. The most serious failure was the breakage of a connecting rod of a helium gas expander. We also had some annoying problems with compressors stopping due to faulty three-phase loss or motor overload detection devices. Since equipment must operate closer to maximum capacity to satisfy demands in the future, such failures will have a more severe impact on cyclotron operation. We continue to try to anticipate problems so remedies are implemented before a weakness in the system leads to time-consuming shutdowns.

Advances in the NSCL Data Acquisition System

R. Fox, R. Au, and A. VanderMolen

Abstract

We report on the advances that have been made in the NSCL data acquisition system over the past year. This report includes a description of our work to develop a network protocol suited to the problem of data acquisition. We also discuss front ends that have been attached to the NSCL data acquisition system via this network protocol. Finally, we discuss the implications of networked data acquisition systems on the back end processing now done on K1200 experiments.

Introduction

Extending the K500 data acquisition system to support the needs of experiments to be run on the K1200 raises several challenges. While we would like to maintain compatibility with the K500 system, there are several extensions required both by the characteristics of the anticipated K1200 experiments and the physical layout of the experimental apparatus.

We anticipate that experiments for the K1200 will involve increasingly complex detection apparatus. This implies an increased parameter multiplicity with associated higher data rates. It also implies more complicated electronics including a variety of programmable trigger and readout modules.

The increased cabling requirements for these devices implies that it is most cost-effective to locate front-end computing close to the experimental areas. This minimizes the number and lengths of analog signal runs. Placing front-end computing near the vault areas also improves the noise immunity of analog signal runs.

The discussion above implies the need for an increasingly intelligent front-end computing system. It also implies that classical point-to-point connections between front end and back end, used in the past, are not usable both because of distance restrictions, and because of a desire to make data available to multiple back-end computers to increase the available on-line analysis power.

Subsequent sections detail the manner in which we have attacked these problems. Sections are organized as follows:

1. **Network protocol developments:** Describes how we have attacked the problems of locating front ends close to the vaults and connections to multiple back ends. We describe the strategy we are pursuing and the status of the projects in this area.
2. **Front-end developments:** Describes how we are attacking the problem of increasing the level of intelligence in the front end processors. We also describe the status of the projects in this area.
3. **Operational aspects:** Describes operational experience with the new system in its current implementation.

Network Protocol Developments

We decided that the most efficient way to satisfy both the front-end to back-end distance requirements, and the desire for multiple back-end processors, would be to connect front ends to the back ends via commercially available local area networking (LAN) hardware. Since ETHERNET interfaces were already present on all NSCL VAX's, both for DECNET and local area VAX clustering (LAVC), we decided to use ETHERNET as a communications medium for the main event data flow.

This strategy has several advantages, namely:

1. Experience exists with ETHERNET cable and subnetwork management.
2. Initially no extra hardware would be needed to make any NSCL VAX into a data acquisition VAX.
3. ETHERNET has very low error rates, and provides isolation via transformer-coupled transceiver taps.

4. Bridges already exist to allow network segmentation to bridge off normal lab ETHERNET from data acquisition subnetworks.
5. There is already substantial software support for this network in the form of DEC software drivers for the ETHERNET interfaces currently available.
6. In VME, UNIBUS and Q-BUS there are already several manufacturers of intelligent ETHERNET adapters which can be used to offload the duties of assured message delivery from the host computers.

On the other hand, the disadvantages are:

1. General-purpose network protocols do not typically provide the data rates required for nuclear data acquisitions. In particular, both DECNET and TCP/IP show rate limits of about 200 kbytes/sec on ETHERNET, which is nearly sufficient bandwidth for a single front-end to back-end link, but not sufficient for multiple links.
2. Good (expensive) diagnostic tools such as ETHERNET protocol analyzers are necessary to track down obscure problems with protocols.
3. Impending FDDI standards are likely to supersede ETHERNET's utility in the data acquisition world.

The considerations above imply that to use ETHERNET would require the development of a specialized communications protocol. In order to support other LAN hardware and software in the future, we wanted to ensure that the front end computer was well insulated from the mechanisms of data delivery. In order to retain compatibility with current data analysis systems, we also wanted to make sure that user level programs were well insulated from the network.

We designed and implemented a streaming protocol based on code which runs in a VME board from Communications Machinery Corporation (CMC) called the ENP-10 (ENP = Ethernet Node Processor). This board is essentially a VME single board computer with a local ETHERNET interface based on the AMD LANCE chip. We programmed the protocol into the ENP-10's MC68010 processor and presented to the VME front end computers a relatively media-independent access mechanism to the protocol. Thus subsequently, the protocol, or other protocols can be ported to other intelligent network interfaces, on possibly other network media, without significant impact on the front end software.

On the VAX, network software was implemented as a new data source to the NSCL ROUTER program. This program already served to isolate on-line analysis programs from the mechanisms of data capture in K500 experiments. This approach also provides automatic compatibility between the K500 and K1200 data acquisition systems and, in fact, K500 on-line data analysis programs have run unaltered under the K1200 system. Placing the ETHERNET data capture part of the protocol into ROUTER also gave us a high-level language view into the debugging of the protocol. This made it possible to design, implement and debug the protocol quickly. The time required to both design and debug the protocol was approximately three months.

We first used ETHERNET to send data between front end and back end computers for the during the first NSCL 4π detector runs last June. It has been used since then for several experiments, including all K1200 experiments to date. The protocol has performed reliably and at good rates under a number of circumstances. Table I summarizes the performance of the protocol.

TABLE I: NSCL network rates.

Receiver CPU	Rate (Kbytes/sec)	Notes
VAXstation 2000	240	1
VAXstation GPX	250	1
VAX-11/750	180	1
VAX 8530	400	2
VAX 8650	400	2,3
VAXstation 3200	300	4
VAXstation 3100	280	4,5

Notes:
1 Rate limited by VAX.
2 Rate limited by ENP-10.
3 Test was across two LANBRIDGE-100's and several hundred meters of broadband cable to a computer in the MSU central site computer lab.
4 This rate is for a real 4π experiment. All other rates are for test data.
5 Preliminary rough measurements only.

In summary, we have created a hardware firmware packaged ETHERNET subsystem. The subsystem employs a protocol developed at NSCL tuned to the problems of data acquisition. This packaged system serves as a system building block which is not only in use in the front end described in the next section, but can serve as a component in future front ends.

In future work we hope to improve the protocol's performance and to extend the functionality of the data acquisition network to include run control. We hope to develop ETHERNET code for Q-bus and UNIBUS-based intelligent ETHERNET processors such as the CMC ENP-40 and ENP-50.

Front-End Developments

This section describes our approach to the need for more intelligence in higher end K1200 experiments. In analyzing the experiments scheduled to run on the K1200, we formed the following impressions which led to the system requirements:

1. Experiments fall into two classes:
   a. Those which involve simple data readout. These can use front end readout routines automatically generated by program generators of the sort used in the K500 system.
   b. Experiments which do either more complicated data acquisition, or must do some device control in addition to data acquisition. These require a system which can easily be modified for special requirements.
2. Users are increasingly interested in getting involved in the tailoring of complex experiments, but they do not have the background to do detailed machine language programming.
3. While cost and performance indicate that CAMAC augmented, by ECL-line, will continue to be the primary readout path for some time, it would not be a good idea to completely close the door on other readout paths such as FASTBUS.
4. While at present, the front ends are based on Motorola MC680x0 architecture CPU's and VME bus, today's rapid developments in processor technology would make it unwise to do front ends in a way which locks NSCL down to that hardware.

An informal and incomplete set of system specifications that emerged were:

1. Use only commercial hardware where possible.
2. Make use of commercial software components where possible.
3. Implement as much of the system as possible in a transportable, high-level language.
4. Provide a clean separation between the sections of the system responsible for data flow and those responsible for interfacing with the event readout devices.
5. Provide high-level language interfaces into all experiment-dependent operations, but provide skeletons which can also call down into machine language routines in case extra speed is required.
6. Provide tools and hooks for user extensions to the system standard functions (e.g. run control command processing of user commands).
7. Provide compatibility with the K500 code generators so simple experiments can be set up as before.

The new front-end system was implemented on a pair of 68020 processors. The system was functionally distributed between an acquisition processor and a control/buffer formatting processor. This model follows closely the separation of duties of the K500 U-2 acquisition front end. Software in both processors is written primarily in the C programming language and runs on Software Components Incorporated pSOS multitasking PROM-based kernel. Independent processes are responsible for data acquisition, interprocessor communication/control, buffer formatting, run control, buffer transmission and data link management. Data is transferred to VAX computers via the ENP-10 board protocol software described in the previous section.

The hardware which makes up the VME system is as follows:

1. Two MicroProject industrial 68020 processors.
2. One Mizar system crate controller.
3. One Mizar 4-port serial interface board
4. Up to 8 Creative Electronics Systems Parallel Branch highway drivers.
5. One ENP-10 ETHERNET processor.
6. Up to 7 CAMAC crates in each of up to 8 parallel branch highways.

The first version of the system was completed in time for the first K1200 experiment. Subsequent experience has shown the software to be very stable. Incremental changes have been made from time to time to improve functionality.

The hardware, however is less stable. The MicroProject processor boards have proven to be unreliable and temperature-sensitive. We are in the process of replacing the processor boards with Ironics model IV-3220 processors. This should not be too difficult because of the clean separation between the hardware-dependent and hardware-independent parts of the software.

The system has since been adapted for use in the NSCL 4π "soccer ball" detector. This adaptation was quite easily done in less than a week. The ease with which this system could be adapted is a good indication that the system software in this system can be regarded as another data acquisition component, a front-end component into which any event readout software can easily be plugged.

Operational Aspects

In the past, NSCL data acquisition front ends were connected to back-end VAXes via various point to point links (CAMAC serial highway or DR-11 compatible). As experimentalists used the networked front ends, it was interesting to see how multiple connectivity changed the way they thought about setting up, to taking and analyzing data on line.

The most notable change was the relegation of the VAX-750's to the role of tape servers, and the use of workstations for on-line analysis. In the point-to-point systems, this was not possible because there could only be one target, and the cost to add a 6250 bpi tape device to a workstation is prohibitive.

To support the use of workstations in on-line and off-line analysis, we produced a version of the NSCL AED display program AEDTSK, which is completely compatible with the original version but runs on 8-plane color VAXstations. This, along with the ETHERNET version of the ROUTER program, allows old 750 acquisition and analysis programs to run unmodified in the on-line environment. In the future, we hope to produce a GKS version of the AEDTSK so that graphics can be displayed on any output device.

With the move to workstations, came the realization that multiple workstations could be used in parallel to analyze data on-line. This makes it possible to provide multiple graphics seats looking at the same data, as well as to use different workstations to perform different analyses. Availability of multiple seats is not just a convenience, but is necessary when electronics is remote from the production running area. This is because it is necessary to look at spectra while adjusting the electronics. Thus, there is often a workstation in the electronics area for the interim K1200 vault as well as one in the area used for production running.

In the case of the 4π system, all electronics is remote-controlled and a local workstation is not necessary. For the 4π however, additional workstations are often used to attempt to do more computationally intensive physics analysis of the data on-line.

Networked acquisition systems allow people besides the experimenters to gain access to the on-line data. In the first K1200 experiment, the data was sent to two workstations for on-line analysis, a 750 for event recording, and a control system workstation to provide additional beamline diagnostics.

Conclusions

NSCL has made major steps forward in the arena of data acquisitions over the past year. We have implemented a networked data acquisition system which has a high level of compatibility with prior NSCL systems. We have maintained the ease of set up and use that has been characteristic of NSCL acquisition systems.

Our network protocol delivers data at rates which are competitive with those provided by labs with similar data acquisition needs. Our new high-level language front end enhances our ability to respond to the specialized needs of our outside and inside users. It has also improved the reliability of the special-purpose front-end codes that we write to satisfy special requirements.

The use of a multisinked acquisition front end, along with workstation software support compatible with old programs, has greatly increased the amount of on-line processing power available to experiments.