OPERATIONAL EXPERIENCE WITH THE BROOKHAVEN ON-LINE DATA FACILITY[/]

S.J. Lindenbaum and S. Ozaki
Brookhaven National Laboratory,
Upton, New York, USA.

A general background description of the Brookhaven On-Line Data Facility (OLDF) and its development since 1964 has been described in a number of prior references [1-3].

The Facility was organized in 1964 as a research support group for on-line experiments utilizing electronic digitized detectors such as counters, counter and spark chamber hodoscopes, etc.

The PDP-6 computer, which was the basic core of the Facility, was, after delivery in January, 1965, debugged to the point where later in 1965 time-sharing was adopted as the operational mode, and multi-user operations have expanded ever since.

Figure 1 shows a typical configuration of users utilizing the PDP-6 system during the previous year or so (prior to the summer of 1969 AGS shutdown). By this time the PDP-6 had been expanded to include 80-96K of 2.0 μsec core (36 bit words), a 450,000 word disc, five high performance magnetic tapes, 8-10 DECtapes, and two printers.* It also included a remote I/O user station which contained a printer and DECtapes.

During this period, the OLDF was simultaneously servicing about five different research user groups. Typically two were on-line via data links simultaneously, with three groups off-line. The disc system was used to swap

---

[/] This work was performed under the auspices of the U. S. Atomic Energy Commission.

* With numerous teletypes and CRT displays distributed over the AGS floor in the user experimental trailers.

users in and out so that typically two large users were in core at any one time. In the swapping mode, each user is assigned a small portion of core continuously ($\sim$ 5-10K) to accept his data link transmissions until they can be transferred to the disc and to do simple necessary monitoring jobs, etc. The bulk of his program is, of course, swapped in and out typically in time intervals which were a few seconds to a minute.

We find that allowing the on-line user program to run locked in core is more convenient for the user (almost instantaneous response time) and allows more efficient use of the processor. However, obviously considerably more memory is required. The system can, of course, accommodate up to 265K of directly addressable memory.

Each user group was provided with its own teletypes and remote scope displays in its experimental trailer. DECtapes and high performance magnetic tape units were assigned to each user. Line printers were shared and a remote I/O station was installed for user convenience and efficiency.

In general, the execution times of a typical computational program are similar to those obtained with the IBM 7094 and about eight to ten times larger than the execution times on a CDC 6600. The time-sharing system allows each user to obtain his own hardware protected core, and monitor protected I/O device allocations, as well as shared I/O devices as desired. The processor is assignable by a scheduling algorithm in rotation for a large number of jobs. A selected number of one-sixtieth of a second time intervals are allocated to a particular job according to the nature of the job and the job load on the system. The processor then works on that job for the assigned time interval and then switches to the next job in the scheduling algorithm automatically and works on that job for the assigned time interval and so on. Thus, provided there are sufficient compute bound jobs in the mix, the processor can be used quite efficiently.

Each user is generally assigned his own teletypes. He can then ask for and obtain (if available) his own hardware protected core, DECtapes (low speed tapes which are primarily used for loading and dumping), his own high performance tapes, and a scope display. The printers are usually shared in rotation, although a user can be assigned a printer. A data link is assigned to each on-line user and allows him to transmit his data to the computer on a priority interrupt basis.

Local Data Handling Equipment and Satellite Computers

The data readout from either counter or (wire) spark chamber hodoscope systems has been predominantly handled by the standard OLDF digital data handling equipment, which is a general purpose system useful for handling mixed information from several different digitized detectors. These units have been provided with 4096 words of 36 bit, 72 bit, and 108 bit word lengths.

The system stacks the bit information from an event received during an AGS pulse sequentially in several consecutive words of its memory. When the buffer store is full, or an AGS pulse ends, after a specified minimum number of words is recorded, the buffer is read out automatically six bits at a time, which are recorded on a high performance tape drive together with a parity bit, and simultaneously a selected fraction of the data is transmitted to the on-line computer over an on-line data link.

A 36 bit unit costs under $15,000 (including the interfaces to experiment and tape drive). The high performance tape units cost about $10,000.

Satellite Computers for Data Handling and Local Diagnostics

With the recent rapid growth of the capacity of small computers, coupled with their rapidly decreasing prices, it has become economically attractive to consider small computers such as PDP-8, PDP-9 and Σ2, etc. for the data handling

function. These computers could then be considered a peripheral processor for the main computer.

The Princeton group has used a PDP-9 in this fashion for data handling and some preliminary local diagnostic work.

The Yale group is similarly using a PDP-8 in this fashion. In both cases, one-way transmission is being utilized from small to large computer. The OLDF is developing two-way transmission devices based on standard interface. These and other groups have found that using small computers at their home institution facilitates preparation for the experiment.

The PDP-15 has been selected as a typical local satellite to be provided to users in the future. A PDP-15/20 system has been ordered and is expected shortly. However, the data link and transmission systems are general enough so that almost any small computer can be interfaced.

The PDP-10 System

The growth in the use of these facilities by the research user community has been such that approximately three quarters of the spark chamber-counter community have or will be utilizing them. Furthermore, the computer requirements of individual user groups have also escalated rapidly as they become familiar with and realize the inherent capabilities of the on-line computer techniques. Thus the OLDF required sufficient expansion to order a PDP-10 system.

The PDP-10 system includes a PDP-10 processor, 16K of 1.0 μsec core and 64K of 1.7 μsec core (Ampex), several high performance magnetic tapes, several DECtapes, and a printer.

The PDP-10 is about twice as fast as the PDP-6 (typical program execution times are about four times CDC 6600 times), but very similar and program compatible so that all PDP-6 programs can be run on the PDP-10. With some minor restrictions on programs, we have been able to maintain the capability of using a program

interchangeably on the PDP-6 or the PDP-10. The PDP-10 and PDP-6 systems have been integrated as shown in Figure 2 so that a common pool of I/O and memory equipment can be utilized as desired for either computer. We typically schedule two on-line users on the PDP-6, several off-line users via disc swapping, and two on-line users on the PDP-10 and have found that our newly expanded facilities are already experiencing the threat of overcommitment. In this regard, it is of interes to note that user program memory core requirements vary from $\approx$ 20-40K. Typically, on-line user requests range between 25% of and 100% of a PDP-6 processor capacity.

The software includes a macro-assembler, an efficient Fortran IV compiler, and a series of debugging, editing, and associated programs. The conversational time-sharing feature allows almost continuous access to the computers simultaneously by a number of programmers. Assistance for user programming problems has been provided by the Facility staff and W.A. Love who has had considerable experience in user research programming problems. At least a sample of the data (and in some cases all of it) is processed on-line. Almost immediate feedback of information is provided by user scope displays and printouts.

In a number of cases, the entire program processing job was done using OLDF computer equipment to complete the off-line processing when spare capacity existed.

The OLDF staff consists of two engineers and five technical specialists. Although we are understaffed due to obvious financial reasons, operational experience has been rather smooth with an overall up-time (during AGS operational periods) of 90%. A major innovation which has kept staff requirements down, is to treat the OLDF as a self-service facility. There are no operators. Each user helps himself in accordance with schedules and equipment and core assignments set-up by the staff. Except for the normal day shift, staff assistance is on an on-ca basis. We find that this self-service system has worked reasonably well and that by mutual negotiation between users and OLDF staff, both efficient operations and

a generally amicable relationship have been developed. This, of course, has only been made possible by the exercise of good judgment and a high degree of cooperative spirit by both users and staff.

In the year or so before shutdown, experiments (completely or partially) performed on-line include the following:

1) A Princeton group's search for CP violation in τ^+ decay. This work was described at the Conference by S. Smith.
2) A study of polarization in $K^+ + p$ by a Yale group.
3) A Columbia group's search for muon pairs in the mass range $\sim 1.5\text{-}6\ (\text{GeV/c})^2$.
4) A BNL group's investigation of $K^\pm \to \pi^\pm \pi^o \gamma$ decays for CP test, and $K\mu_3$ decay and polarization.
5) Successful tests of the Double Vee Magnetic Spectrometer system by a BNL group.

Figure 2 shows the user load and mode of operation in the present and projected near future AGS experimental operations.

REFERENCES

1. S. J. Lindenbaum. Digitized spark chamber and on-line computer systems for double vertex events. Proceedings of the International Symposium on Nuclear Electronics, Versailles, September 10-13, 1968, Volume III, pp. 50-1 - 50-25, 1968.

2. S. J. Lindenbaum. Large digitized spark chamber spectrometer and on-line computer systems. Proceedings of the International Conference on Advanced Data Processing for Bubble and Spark Chambers, Argonne, October 28-30, 1968, Edited by R. J. Royston, ANL 7515, pp. 202-225, 1969.

3. W. A. Love. Brookhaven on-line data facility. Proceedings Skytop Conference on Computer Systems in Experimental Nuclear Physics, CONF-690301, pp. 124-38, U.S. Atomic Energy Commission, 1969.

FIGURE CAPTIONS

Figure 1: The general typical use of the PDP-6 facility during the prior AGS schedule period until the May 1969 shutdown. Two groups were put on-line at a time via the patch panel and two or more groups prepare programs, etc. off-line simultaneously. Memory protection and relocation hardware make this time-sharing operation, which has been in use for more than two years, efficient and routine. The two trailers shown are 40' long x 10' wide truck trailers with air conditioning. They are located just outside the east experimental building of the AGS.

Figure 2: The arrangement for the combined PDP-6 and PDP-10 facility. Each computer system is shown in a 40' long x 10' wide trailer with the average minimum memory and I/O equipment usually committed to its use. The trailer in the middle contains a common pool of memory and I/O equipment which can be easily switched to either computer as desired. Even the memory and I/O in the outside trailers can be switched to either computer if desired. The solid line boxes indicate equipment in use now. The dotted line boxes indicate the major additional items expected to be incorporated shortly.

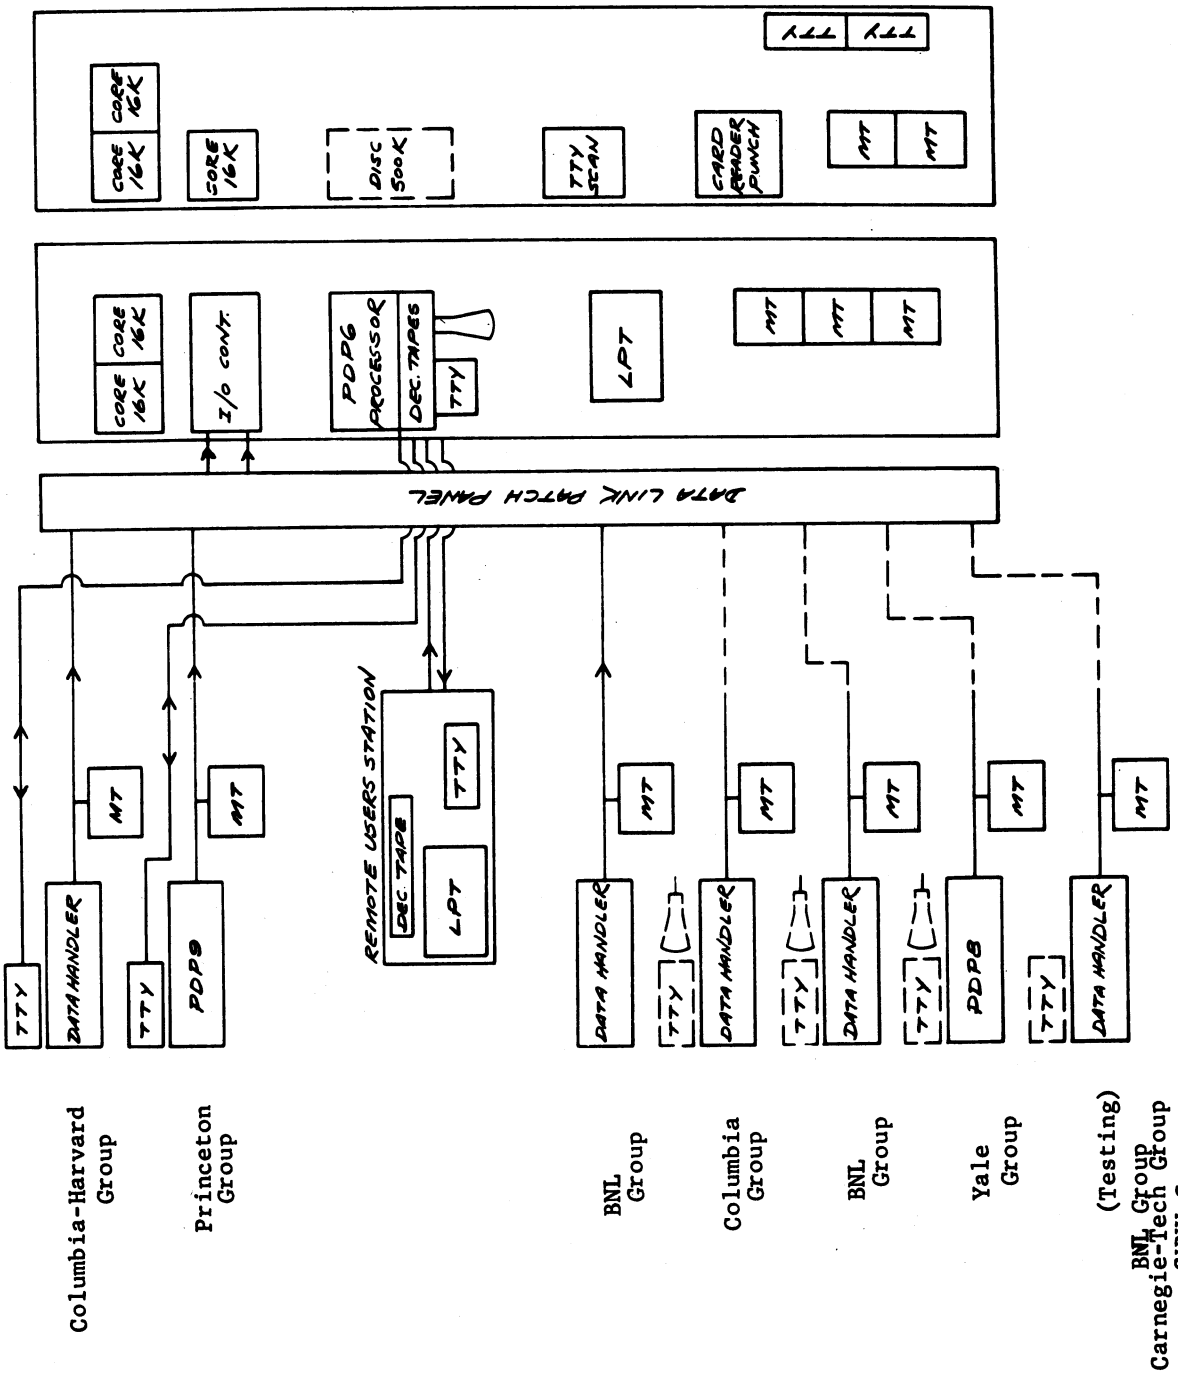
CORE 16K
CORE 16K
CORE 16K
CORE 16K
CORE 16K
DISC 500K
TTY SCAN
CARD READER PUNCH
TTY
TTY
MT
MT
CORE 16K
CORE 16K
I/O CONT.
PDP6 PROCESSOR
DEC. TAPES
TTY
LPT
MT
MT
MT
DATA LINK PATCH PANEL
REMOTE USERS STATION
DEC. TAPE
TTY
LPT
TTY
DATA HANDLER
PDP8
MT
Columbia-Harvard Group
Princeton Group
BNL Group
Columbia Group
BNL Group
Yale Group
(Testing)
BNL Group
Carnegie-Tech Group
SUNY Group

FIGURE 1

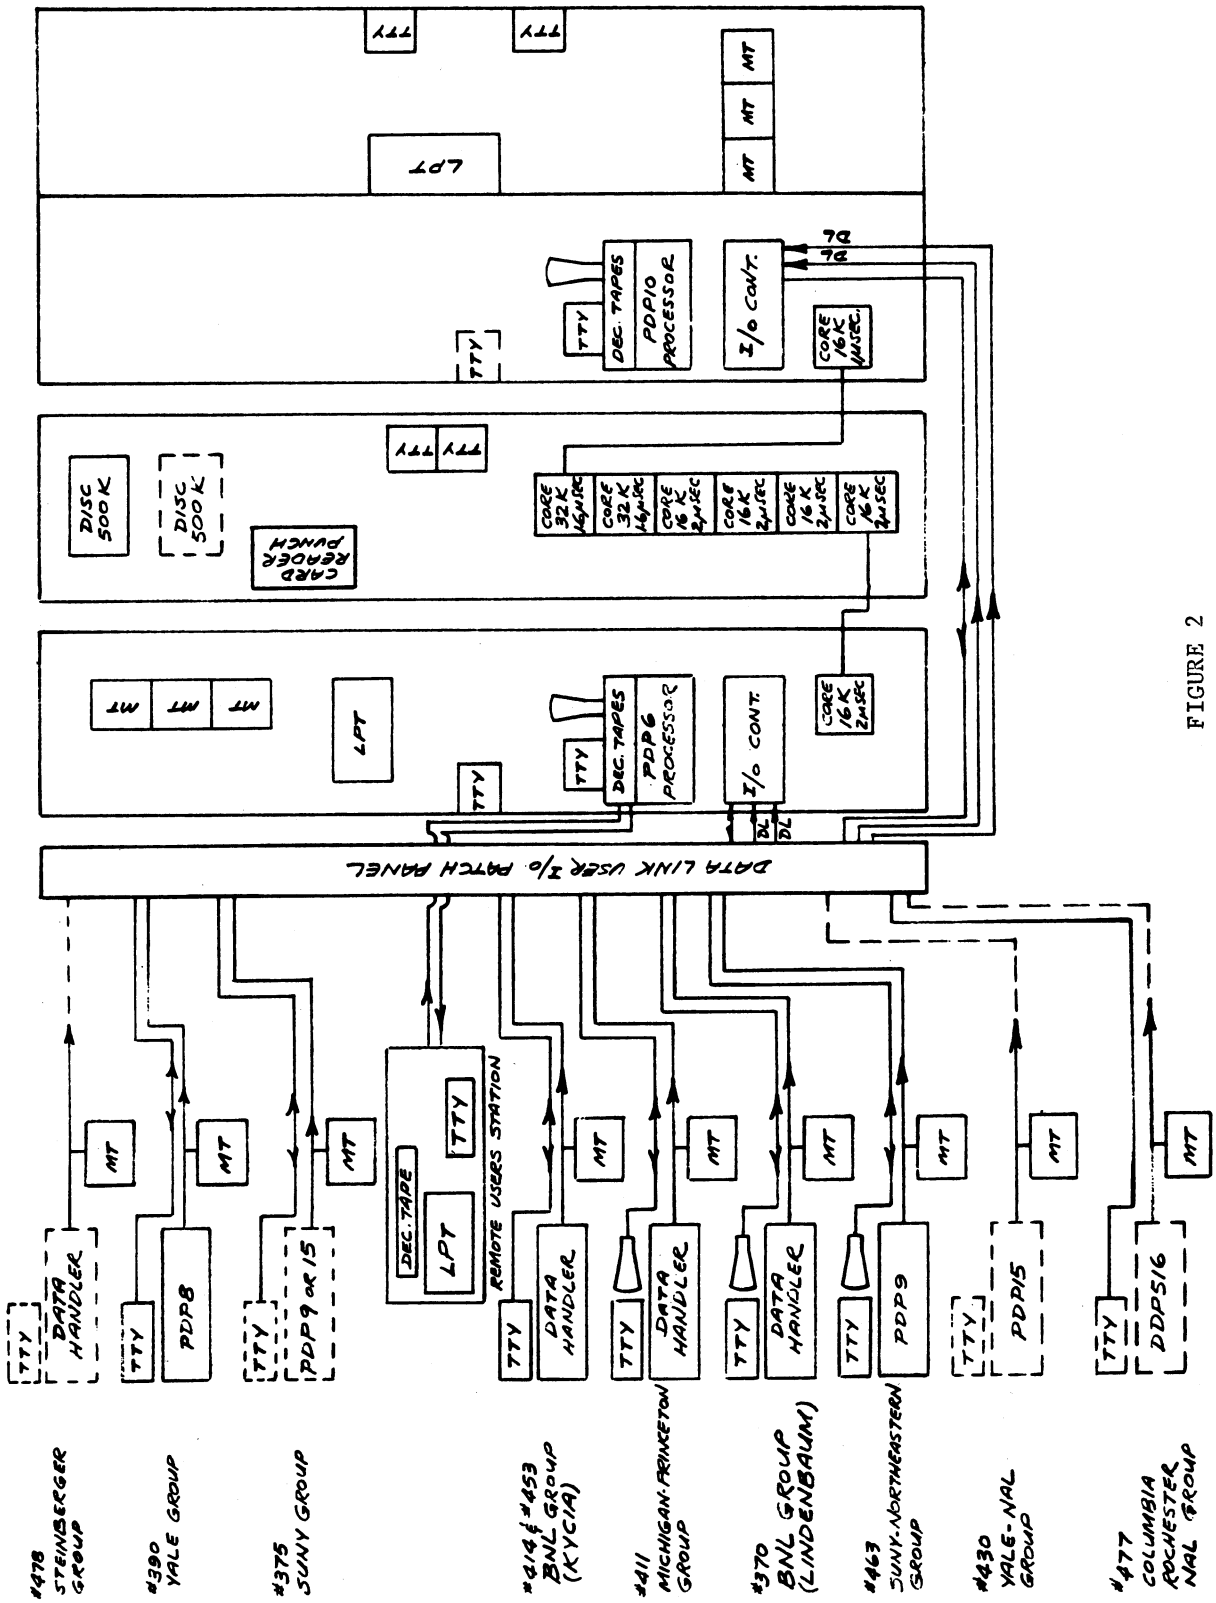
#478 STEINBERGER GROUP
#390 YALE GROUP
#375 SUNY GROUP
#414 & #453 BNL GROUP (KYCIA)
#411 MICHIGAN-PRINCETON GROUP
#370 BNL GROUP (LINDENBAUM)
#463 SUNY-NORTHEASTERN GROUP
#430 YALE-NAL GROUP
#477 COLUMBIA ROCHESTER NAL GROUP
DATA HANDLER
PDP8
PDP9 OR 15
DEC. TAPE
LPT
TTY
REMOTE USERS STATION
PDP9
PDP15
DDP516
MT
DATA LINK USER I/O PATCH PANEL
DEC. TAPES
PDP6 PROCESSOR
I/O CONT.
CORE 16K 2μSEC
PDP10 PROCESSOR
CORE 32K 1.6μSEC
CORE 16K 2μSEC
DISC 500K
CARD READER PUNCH
777

FIGURE 2