WIRE-SPARK-CHAMBER SYSTEM WITH ON-LINE COMPUTERS*

W.T. Ford†, P.A. Piroué††, R.S. Remmel,
A.J.S. Smith and P.A. Souder.

Department of Physics, Joseph Henry Laboratories,
Princeton University, Princeton, N.J. 08540, U.S.A.

1. INTRODUCTION

This paper describes an on-line wire-spark-chamber system used at the Brookhaven A.G.S. to measure the τ decay mode of K^+ and K^- mesons ($\tau^\pm \equiv K^\pm \rightarrow \pi^\pm \pi^\pm \pi^\mp$). Slightly more than 3 million good events were detected. The data handling system, with a PDP-9 computer interfaced to the PDP-6 computer[1] of the Brookhaven On-Line Data Facility (OLDF) may be of interest for the following reasons:

i.) Substantially the same method as that described here is used at present and will be used in the foreseeable future to interface on-line small computers to the OLDF.

ii.) This was the first experiment at the A.G.S. to use both small and large computers on line, and also one of two initial experiments to use the OLDF's job-swapping disk system, with which two on-line users shared both PDP-6 core and processor time with up to 4 off-line users.

---

*Work supported by U.S. Atomic Energy Commission Contract AT(30-1)-2137 and Office of Naval Research Contract No. N0014-67-A-0151-0001.

†Permanent address: Synchrotron Laboratory, California Institute of Technology, Pasadena, California 91109

††Present address: CERN, 1211 Geneva 23, Switzerland

iii.) Our design might serve as a prototype or basis of comparison for outside user groups planning on-line experiments, who face as we did the problems of constructing and testing much of their apparatus and software before arriving at the accelerator.

In what follows we shall outline the experiment and the physics behind it, and then discuss the acquisition and processing of the data.

2. THE EXPERIMENT

2.1 Motivation

All existing data on τ decay[2,3] are consistent with the linear approximation to the τ-decay matrix element:

$$|M|^2 \, d\,X \, d\,Y = \text{const} \, (1 + a_\tau Y) \, d\,X \, d\,Y$$

where a_τ is a constant, and X, Y are the usual Dalitz variables. A high-statistics measurement of τ decays can give new information as to the existence of higher order terms in the matrix element, e.g. X^2, Y^2, X^2Y terms, as well a more precise value for the slope parameter a_τ.

Such knowledge is interesting primarily for the following reasons:

i.) Test of CP invariance.

Any difference between the Dalitz plots of τ^+ and τ^- would indicate a CP violation outside the neutral kaon system and hence one which could not occur via the "superweak" interaction.[4]

ii.) Test of $\Delta I = 1/2$ rule.

The comparison of the parameter a_τ for τ decay with the corresponding parameter $a_{\tau'}$ for τ' decay ($K^{\pm} \to \pi^{\pm} \pi^{0} \pi^{0}$) tests the $\Delta I = 1/2$ rule, which predicts[5,6] that

$$a_{\tau'}/a_{\tau} = -2.0$$ neglecting electromagnetic corrections.

We hope to reduce the uncertainty in a_{τ} by a factor of ~ 4.

2.2 Experimental Arrangement

The experimental arrangement is shown in Fig. 1. K^+ or K^- mesons in a charged 3.0 GeV/c separated beam were identified by a differential gas Cerenkov counter and 3 scintillation counters (only the last beam counter is shown in the Figure). Kaons then decaying into 3 charged pions were detected with less than 2% background by a 3-fold coincidence in an 8-counter hodoscope octagonally symmetric about the beam. For 58% of such coincidences, the "unlike" pion (opposite charge to the beam) was detected by a 16-counter array behind the wide-aperture spectrometer magnet D_4; wire spark chambers were then triggered with 20 - millisecond dead-time. (All "like" pions and beam particles missed the rear hodoscope and spark chambers completely). The 5 chambers (each with X and Y planes) ahead of D_4 measured all 3 pion directions. Each chamber had > 95% 3-spark efficiency/spark with $\sim$ 25% probability for an extra spark. Two of the chambers were rotated 45° w.r.t. the other three to resolve ambiguities. The six chambers behind D_4 served to measure the unlike pion momentum.

The spark coordinates were digitized via a magnetostrictive-readout and scaler system purchased from Science Accessories Corp. Each event consisted of 90 18-bit words, including hodoscope information. Typically 10-15 events were accepted in each 400-msec A.G.S. burst.

3. DATA ACQUISITION AND ANALYSIS

3.1 Principles of Operation

The data flow is shown in Fig. 2. For on-line operation, digitized information from the experiment was fed to a PDP-9 computer, programmed in assembly language, which buffered the data and reinitialized the spark chamber read-out system. Between A.G.S. bursts the raw data in the buffer was dumped onto magnetic tape, and simultaneously transferred via an interface to the PDP-6, where complete reconstruction of a substantial fraction of the events took place. Summarized results from the PDP-6 concerning the apparatus were printed on a remote teletype in the experimental trailer. More complete results of the on-line analysis were put on magnetic tape and a line printer.

Off-line analysis was facilitated by the fact that the reconstruction program also accepted input either from the data tapes, or from Monte-Carlo-generated data. Written in Fortran, this program was source-deck-compatible among the three installations at which it was used. This made it possible for us to develop the reconstruction programs on the

Princeton 7094 or 360 computer before moving to Brookhaven, and to perform off-line analysis at the B.N.L. 6600 computer during the run. To ease the reading of tapes on these various machines, the tape-record length was fixed at 4080 6-bit characters, an integral number of words for the 7094, PDP-9, PDP-6, CDC-6600, and IBM 360.

3.2 The PDP-9 System and its Interfaces

i.) PDP-9 Performance.

The PDP-9 configuration is now summarized:

- memory: 8192 words, 18 bits, 1 μsec cycle time.
- Extended arithmetic element: hardware multiply, divide, normalize, shift instructions (fixed point only)
- Magnetic tapes: 1 TU-20 unit (36,000 characters/sec at 800 b.p.i.)
  2 TU-55 DEC tapes for storing programs.
- Data Channel Hardware: Steals 3 central-processor cycles/word to transfer data between external devices and memory.
- Program interrupt: one level of priority.

As well as controlling the data flow, the PDP-9 monitored the performance of the spark chambers and hodoscopes. For the front chambers the sparks were fitted to tracks; the tracks were then tested to see how many came from a single vertex. The residuals to the track fits were computed for the 10 wire planes (2 planes/chamber) and stored as histograms in a display scope described below. Also computed and displayed were the raw spark-frequency distributions for all 22 planes and two-dimensional illumination of any two chambers selected by the operator. These displays were used to align the chambers by software, and to expose malfunctions such as edge-breakdown, inefficiencies, etc. Serious error conditions caused error messages to be printed on the teletype. The computing speed (~ 7 events/sec) was more than adequate to process all events, even at the maximum event rate of ~ 15 events/burst.

The PDP-9 system was designed so that it could operate completely independently of the PDP-6; because of this we could take meaningful data even when the PDP-6 was not working. More importantly this independence made it possible to test the PDP-9 system and spark chambers before moving to Brookhaven. This latter feature seems to be essential for outside users, if they are to get a complex apparatus working under the pressure of accelerator schedules.

ii.) Display Scope

Our display scope consisted of a Nuclear Data 4096-word, 18-bit pulse-height analyzer memory unit interfaced as a multiscaler. The hardware-wired programs of this unit displayed the contents of the memory on a Tektronix 503 oscilloscope. To the PDP-9 this device appeared just like a teletype, raising its flag when ready to accept another word. Aside from the obvious saving of 4096 core locations, this device also freed the PDP-9 from the processor-time-consuming display function.

iii.) Interface to PDP-6

An extremely simple interface was possible because the PDP-6 data-link accepted 6-bit characters, packing them 6 characters per 36-bit word, and was handled by the PDP-6 similarly to a magnetic tape unit. Hence we simply tapped the 7 data signals (6 bits + parity bit) from the write circuits of our TU-20 tape unit, amplified them, and sent them to the Data-Link. The data transfer rate was thus 36,000 characters/sec) the speed of the TU-20 tape unit. Only three additional lines were needed:

- A D.C. reference signal to eliminate common-mode noise.
- A line from the PDP-9 to the Data-Link requesting service.
- A busy/not-busy line from the Data-Link to the PDP-9.

Note: The Data-Link was designed originally to interface only Data-Handlers(4096-word storage memories) to the PDP-6. However, the simplicity of interfacing small computers such as ours to it has led to its use in all scheduled experiments with small computers on-line to the PDP-6.

3.3 On-Line Use of the PDP-6

i.) Purpose of On-Line Reconstruction.

To operate consistently, the spark-chamber system had to be finely tuned. Chamber performance could severely deteriorate in a few minutes, in ways not discernable other than by complete on-line reconstruction of a large fraction of the data. For example, the rate of spurious sparks could only be found after all tracks had been fitted. Similarly, to determine the properties of the last 6 chambers, we had to reconstruct particle trajectories through the spectrometer magnet. On the basis of such information, adjustments of chamber parameters were made when necessary, and soon enough to prevent any accumulation of useless data.

Needless to say, a large computer with fast floating point hardware was needed for this task, as well as an efficient Fortran compiler. The PDP-6 of the OLDF fulfilled these requirements most satisfactorily.

ii.) The Job-Swapping System.

This system has been described in detail elsewhere,[1] but a few comments from a user's point of view are perhaps in order. (The last half of our data were taken using the swapping system, it having been installed during our experiment.)

The OLDF hardware configuration as of January 1969 is shown in Fig. 3. The disk had 500k 36-bit-word storage, 17 msec access time, and 13 μsec/word transfer time. The solid lines denote experiments in progress or completed; the dotted lines indicate off-line users testing their programs for "upcoming" experiments. We shared on-line use with the Columbia (Lederman) group, both groups having ~ 25k compute-bound programs.

The Columbia program remained in core, while ours was swapped with off-line users every 10 sec or so. The sharing of CPU time was adjusted independently to keep the computer busy. We usually received ~ 50% of the CPU time, as did Columbia. (Off-line users usually required negligible CPU time.) Because there was always a compute-bound job in core, CPU efficiency was virtually as good as when both on-line jobs resided in core. Typically we analysed ~ 30% of our data on-line, the analysis time per event being ~.3 second. System reliability, though marginal at first, improved rapidly, till by the end of our experiment system failures were not a problem.

4. CONCLUSION

The system just described is not by any means the most efficient from a strictly instrumentational point of view. However, its flexibility and simplicity made it extremely practical, and possibly the only system with which our relatively small group could have assembled the experiment in a reasonable time. The use of Fortran for the reconstruction programs meant we did not have to become programming specialists.

PDP-9 assembly language was relatively straightforward, and very efficient in performing the logical operation for which it was used. The PDP-9 was built into the apparatus from the beginning; by the time we got to the A.G.S., it was a familiar piece of equipment. Finally, and perhaps most important, the many components of apparatus and software could be constructed and tested independently, in our own laboratories and computer center, and then carried over to the final on-line configuration with a minimum duplication of effort.

5. ACKNOWLEDGEMENT

The authors wish to thank the many people who made this work possible. Dr. S. J. Lindenbaum, Dr. S. Ozaki, Dr. W. Love, and the staff of the OLDF worked very hard to help us. Mr. J. Gould and Mr. A. Friedman gave essential advice and support concerning the hardware and software links to the PDP-6, and were most cooperative in solving problems as they arose. At Princeton Mr. K. Wright, Mr. Fred Schwarz, Mr. Howard Edwards and their staff built much of the apparatus. Mr. Edward Card fabricated much of the digital electronics. It is with special gratitude that we acknowledge Mr. H. Meusel, who sustained a serious and unfortunate injury while working on our equipment. One of us (PAP) wishes to thank Professor B. P. Gregory and Professor G. Cocconi for their hospitality at CERN.

This work made extensive use of the Princeton University Computer Center, sponsored in part by the National Science Foundation.

REFERENCES

1) For a description of the On-Line Data Facility see the paper submitted to this conference by S. J. Lindenbaum and S. Ozaki.

2) T. S. Mast et al., Phys. Rev. 183, 1200 (1969).

3) J. Grauman et al., Phys. Rev. Letters 23, 737 (1969).

4) L. Wolfenstein, "Ettore Majorana" Summer School Notes, 1968 (unpublished).

5) S. Weinberg, Phys. Rev. Letters 4, 87 (1960).

6) R. F. Sawyer and K. C. Wali, N.C. 17, 938 (1960).

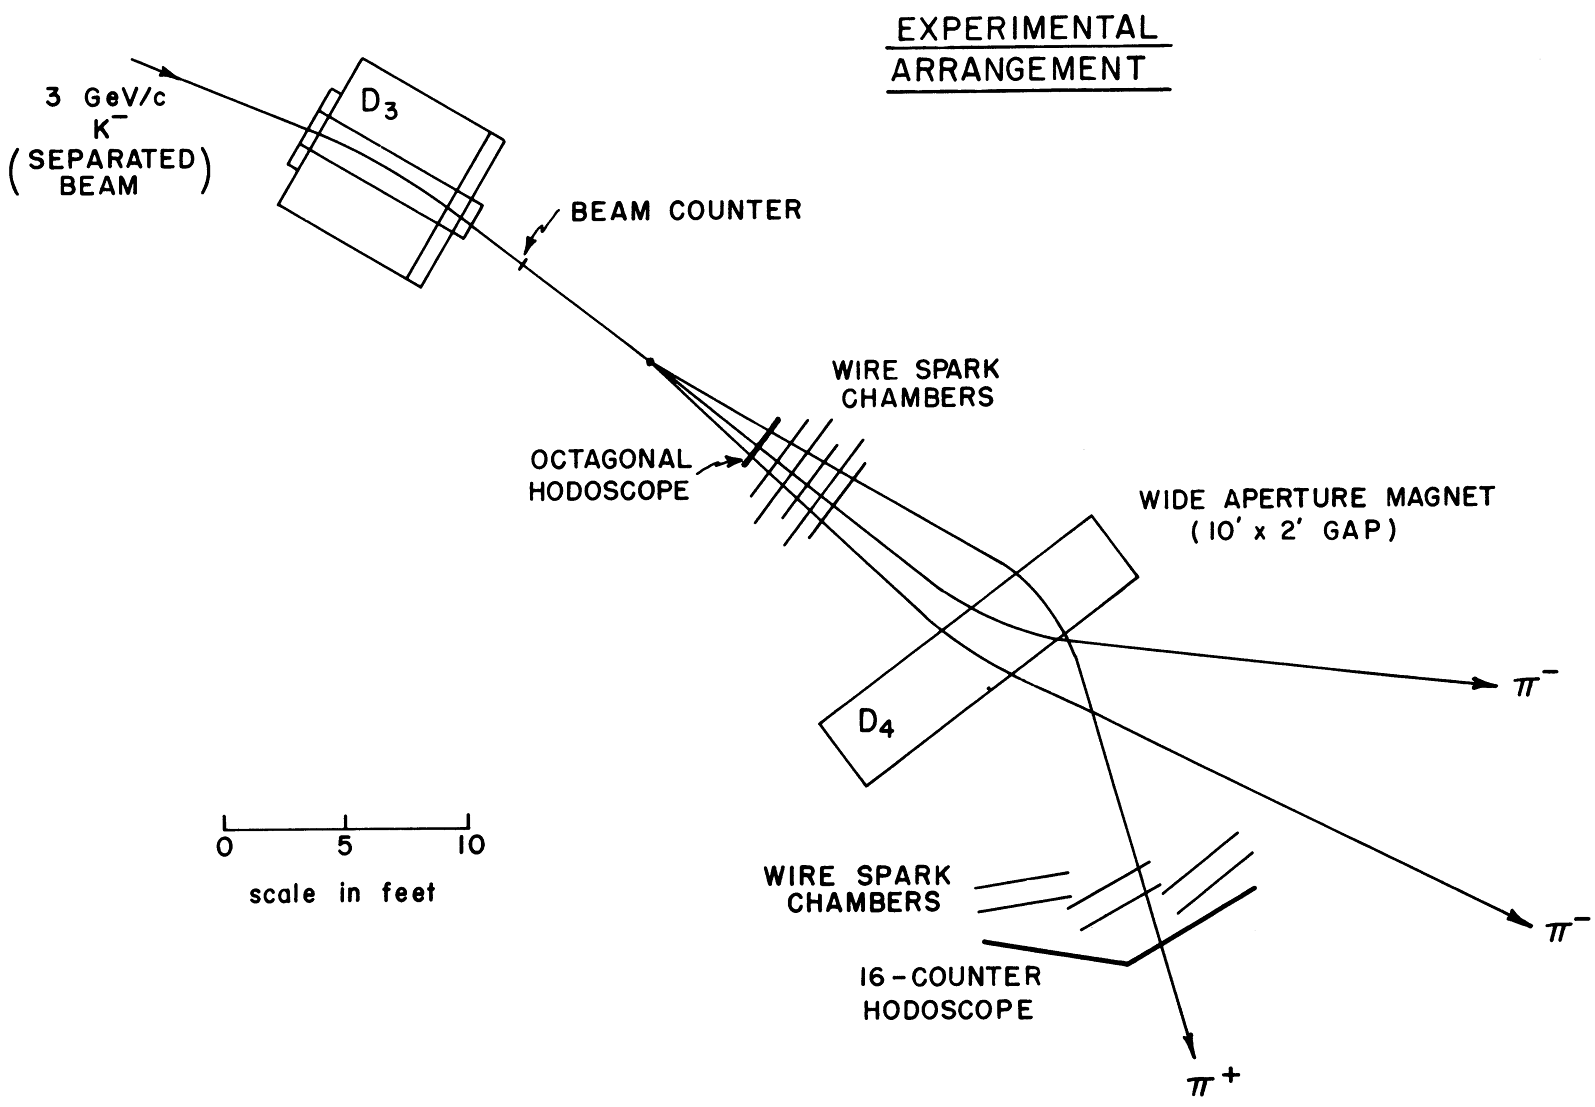

Fig. 1. Experimental arrangement

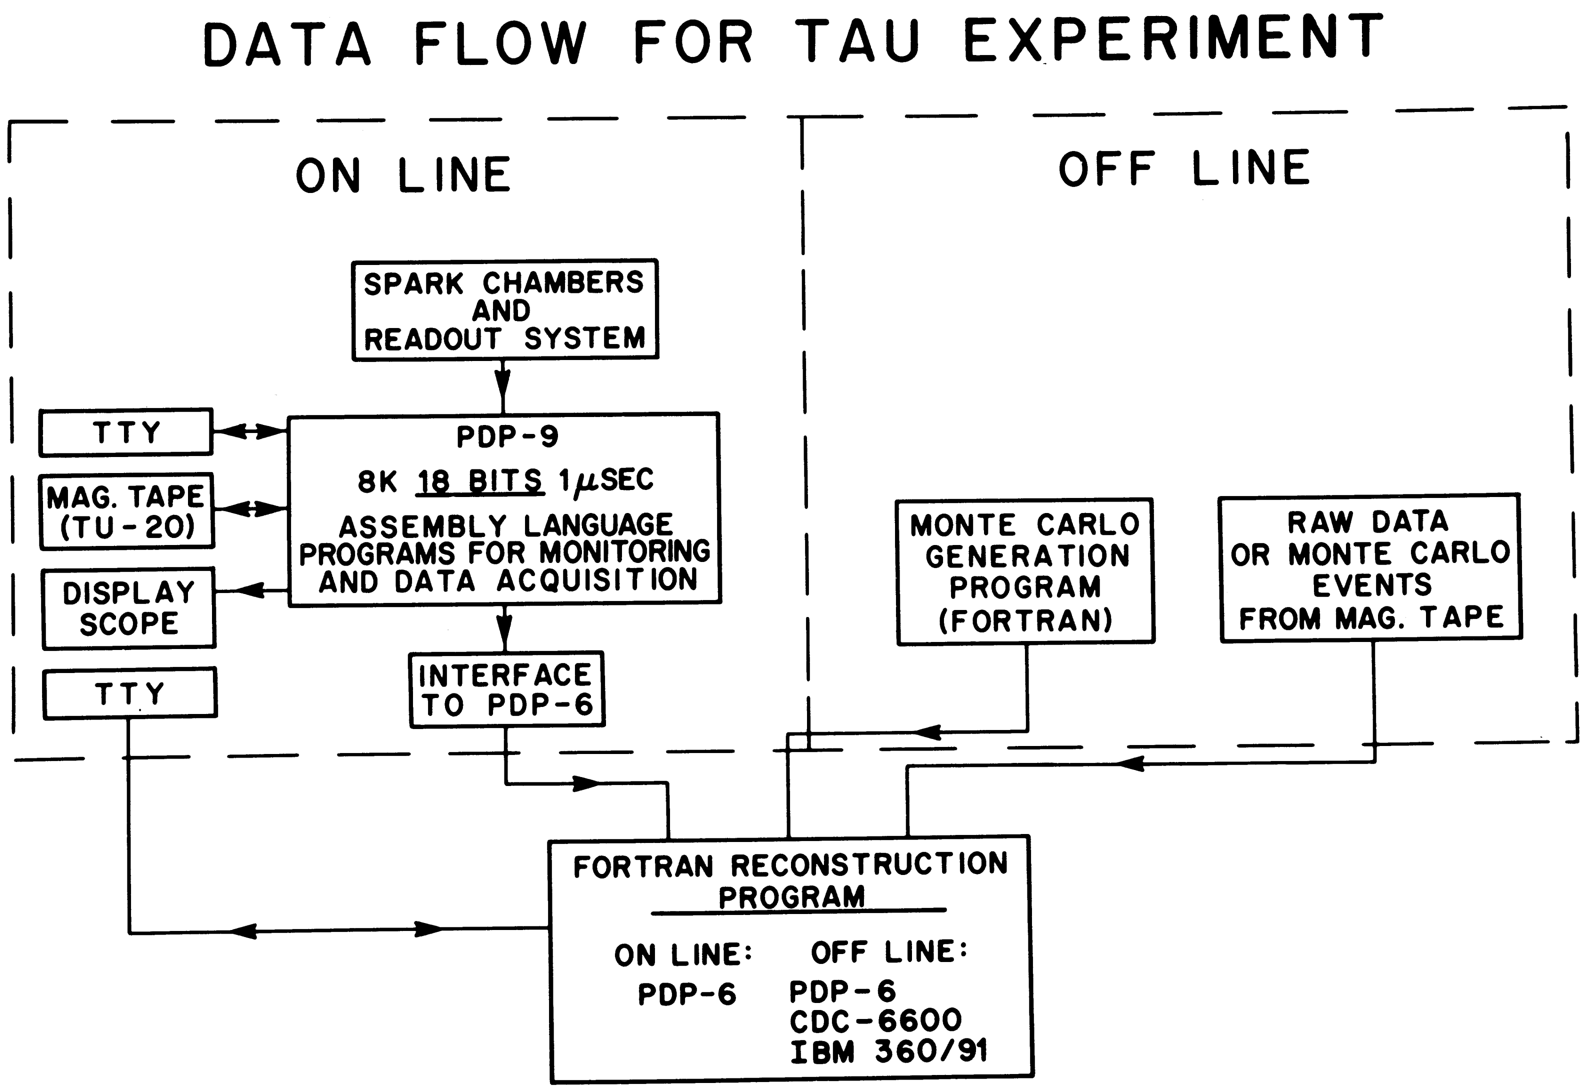

Fig. 2. Data flow. A detailed explanation is given in the text. The part labelled on-line was in use during the taking of data. The same reconstruction program was used on-line and off-line.

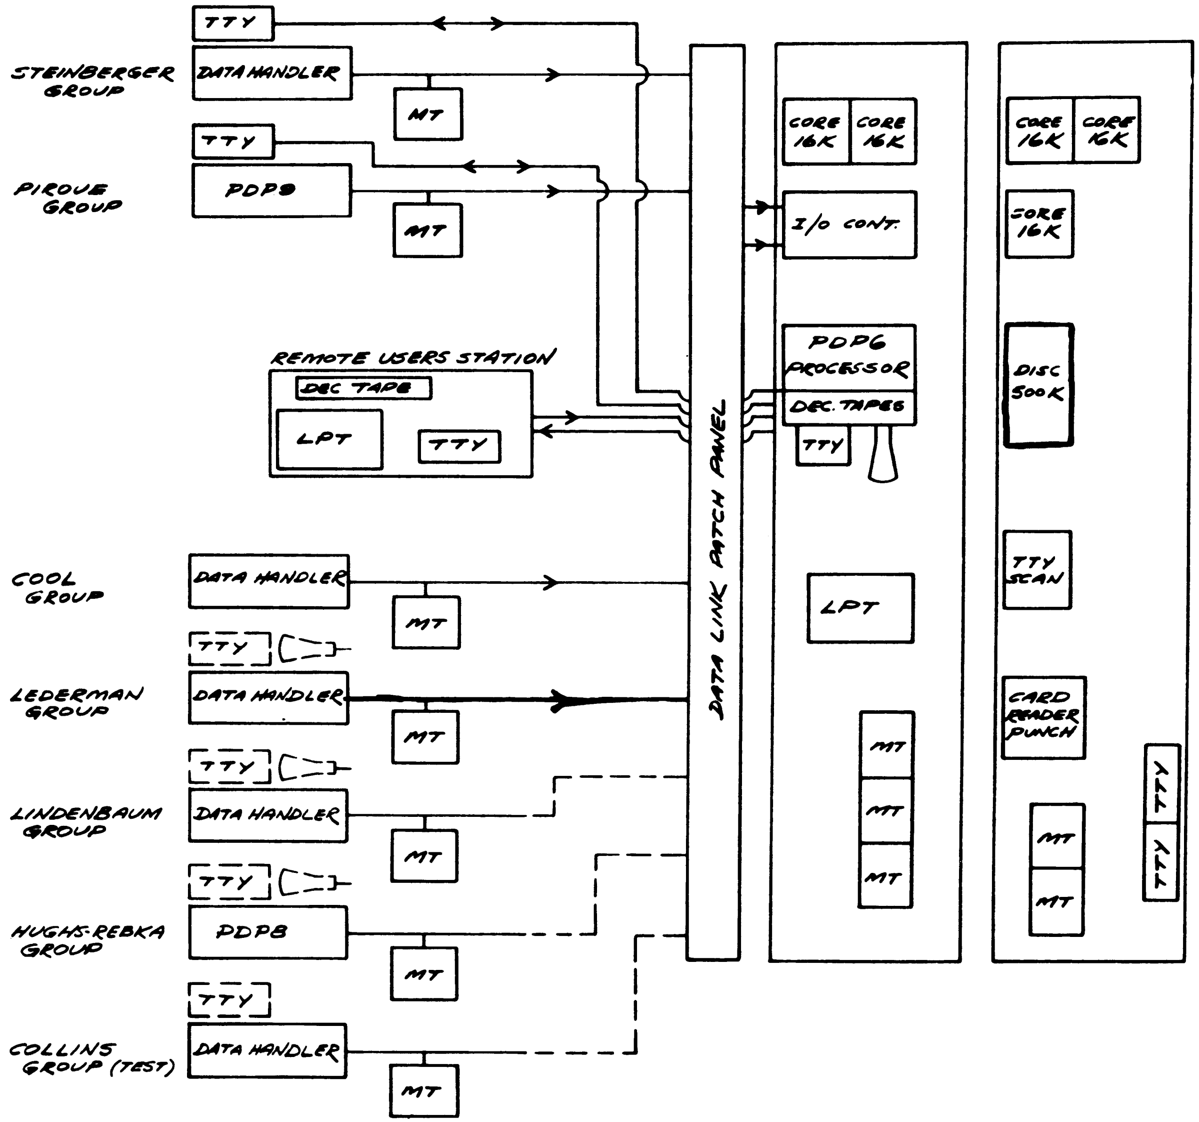

Fig. 3. The On-Line Data Facility as of January, 1969. The PDP-6 is housed in two trailers depicted at the right of the figure. The experiments on the floor are represented at the left by solid lines. Dotted lines are used to show those experiments preparing programs off-line.

DISCUSSION

H. FAISSNER *(Aachen):* How many accidental tracks did you typically register in your chamber?

A.J.S. SMITH: Accidental tracks were five percent, I would say. It was a separated beam, the pion-to-K ratio being about ten to one.

H. FAISSNER *(Aachen):* Did you encounter any ambiguities in connecting tracks before and after the magnet?

A.J.S. SMITH: There was a small amount, we were greatly helped by one feature which is rather amusing. I mentioned that the magnetic field shape of that magnet is terrible, but it provides a significant focusing and large vertical displacements, so by matching up the vertical and horizontal position of the track there were not many ambiguities, less than one percent.

R. BAIRSTOW *(RHEL):* What was the resolution for two sparks?

A.J.S. SMTIH: We managed to get a separation between two sparks of about six to seven clock counts, one clock count corresponded to something between one and two millimetres. The wire separation is 52 wires per inch.

R. BAIRSTOW *(RHEL):* What was the effect of the magnetic field?

A.J.S. SMITH: The magnetic field was an extremely serious problem to start with, and it was solved simply with brute force by sticking enough iron between the big magnet and our chambers. We put two inches of iron along all the magneto-strictive delay lines.

P. VILLEMOES *(CERN):* What was the data rate from the experiment?

A.J.S. SMITH: In any given AGS pulse we took about 15 events, and each event had about 90 8-bit words in it.

D. RUST *(Argonne):* Did you have to re-analyse all events off-line even though 20% to 40% had been analysed on-line?

A.J.S. SMITH: Unfortunately, yes. This is partly due to the Monte Carlo preparation of the events not being sufficient -- one does not know the background and other problems arise during the run. I think that especially for a small group there is not much sense in making the on-line program so elaborate that one can hope to avoid processing these events again. One has to. One should regard the on-line processing more in the spirit that it tells you what is going on.

R. ROSNER *(RHEL):* What sort of output came back to the experimenter from the PDP-6 and how often did such output occur?

A.J.S. SMITH: The main information we got back from the PDP-6 on-line was the efficiency of each chamber and the amount of extra sparking in each chamber. This came back every ten minutes, unless there was something wrong, for example it found that the efficiency of one of the chambers dropped below 93 per cent. In this or in similar cases it would print an error message. We detected some extra sparking in this way, which we could overcome by adding alcohol to the gas.

K.M. SMITH *(Glasgow):* What resolution was achieved in reconstruction and how long did it take to do it?

A.J.S. SMITH: On the PDP-9, we processed about seven events per second. The reconstruction of the vertex was accurate to about half an inch. On the PDP-6, it was slightly better. Including the reconstruction through the magnet, it took about 300 msec per event.

B. POWELL *(CERN):* Did you feel that in general it would be useful to have more computer capacity available on-line than that provided at BNL?

A.J.S. SMITH: I would say that what we had was sufficient, but of course the experiment was not very complicated. For a larger experiment it would not have been sufficient. We needed the results of the on-line analysis of about 30 per cent of all events, and this would probably hold for other experiments as well. So, if the reconstruction takes much more time, you need more on-line capacity.

D. WEBSDALE *(CERN):* Which principal factor limited your event collection rate to about 15 per burst?

A.J.S. SMITH: The main limit was set by the dead time of the chambers. If we made it much less than 20 msec, they would not fire reliably, since we did not have a pulsed clearing field circuit. But the number of useful events per burst was not very much higher.