

Durham E-Theses

Bubble chamber film: the development of an online control system

C. Thompson

How to cite:

Thompson, C. (1971) Bubble chamber film: the development of an online control system. Masters thesis, Durham University.

Use policy

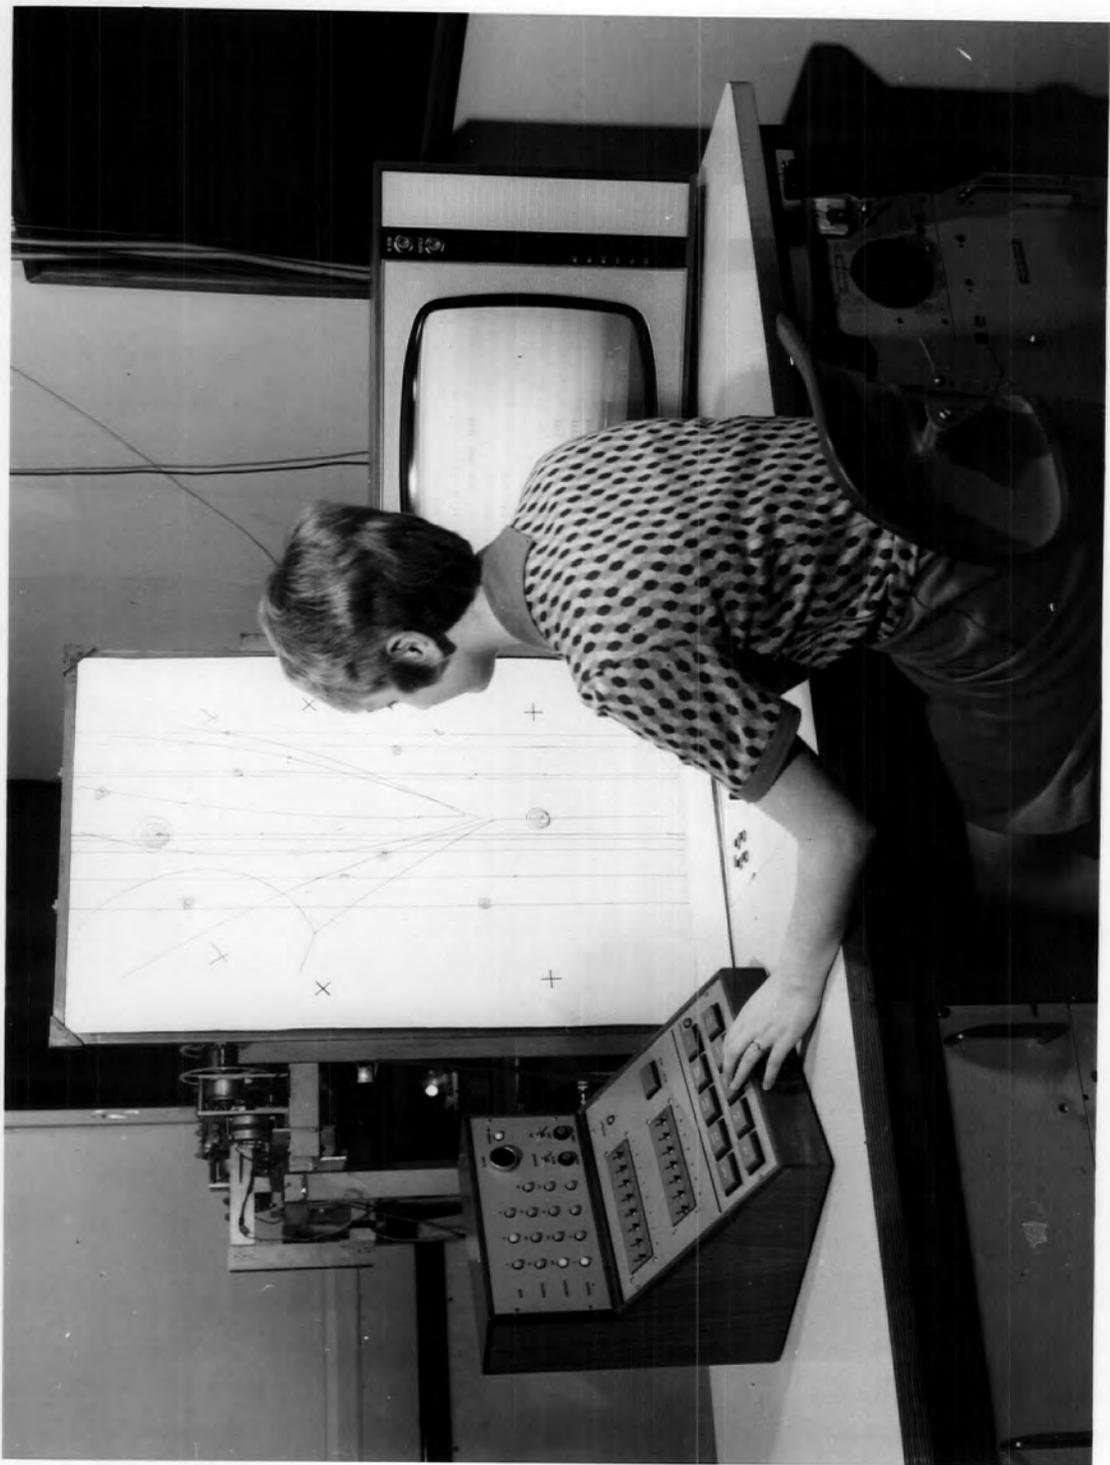
The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a <https://etheses.durham.ac.uk/id/eprint/10254/> is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full Durham E-Theses policy](#) for further details.

Frontispiece - Operator's Console



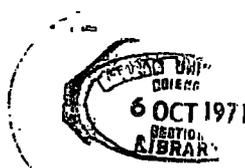
BUBBLE CHAMBER FILM MEASUREMENTS:
THE DEVELOPMENT OF AN ON-LINE CONTROL SYSTEM

by

C. Thompson, B.Sc.

A thesis submitted to the
University of Durham
for the degree of Master of Science

June 1971



ABSTRACT

An on-line system has been developed for the measuring of events on bubble chamber film. The system is controlled by an IBM 1130 computer with 8K words of core, and a core cycle time of 3.6 μ sec. Up to six measuring machines may be operated simultaneously, and the system will also handle the collection of data from a core store which is part of a cosmic ray experiment.

At present four machines are operating on-line. These had previously been used for off-line measuring when the data were punched onto paper tape.

A comparison between the off-line system and the on-line system of measuring reveals that the development of the latter system has resulted in an increase in the measuring rate of approximately 50%, and a similar increase in the percentage of events which are successfully reconstructed by the geometry programme THRESH. As a result, the number of events that are successfully measured per machine per hour has more than doubled.

PREFACE

This thesis is an account of the work carried out by the author while she was working under the supervision of Dr. J. V. Major as a Research Assistant with the Bubble Chamber Film Analysis Group in the University of Durham. When the author started to work in Durham, in November 1967, the Group was using a paper tape measuring system, and was processing the data on an Elliott 803 computer in Durham and an IBM 7044 computer in Glasgow. Within the following 18 months substantial changes took place, and the author was primarily responsible for writing and testing many of the new programmes that were required.

Firstly, the Elliott 803 was replaced by an IBM 1130 computer, which was later linked to an IBM 360/67 computer in Newcastle as part of the Northumbrian Universities Multiple Access Computer (NUMAC). The paper tape output from the measuring machines was then edited and sorted on the IBM 1130, and its card output was sent to Glasgow. Later when the required programmes were working on the IBM 360/67 the cards were taken to Newcastle. The new programmes which were required to edit and sort the data on the IBM 1130 computer were written and tested by the author.

The second change was the development of an on-line system of measuring to replace the paper tape system. The on-line system, which will accommodate up to six measuring machines, is controlled by an IBM 1130 computer with 8K words of core. The data from this system are output onto disk, and are later input to NUMAC for further processing via the Durham terminal computer. The author was involved in designing the

on-line system, and was entirely responsible for writing and testing all the programmes that were required to monitor the operation of the measuring machines, to check the measurements as they were made, and to edit and sort the data before they were input to NUMAC.

Similar on-line systems were already being developed at CERN¹, Westfield College², the University of Liverpool³, the Rutherford Laboratories⁴, and the University of Oxford⁵, and others have also been developed at the Universities of Edinburgh⁶ and Turin⁷. These systems, like the Durham system, are controlled by small computers such as the IBM 1130 or DEC PDP8 with 8K or 16K words of core.

In the majority of cases the operators of the measuring machines communicate with the computer by means of interrupts. The alternative method which is used at the Rutherford Laboratories, requires the computer to interrogate each machine in turn to see if it has data available for reading in.

Most of these systems are operator dominated. That is, each operator is free to make the measurements for each event in any order, and she must therefore indicate at each stage what she is measuring. The Durham system on the other hand is computer dominated. That is, the computer controls the order of measuring by calling for each of the measurements in turn, and the operator responds by making the required measurements.

A full description of the Durham on-line system is given in this thesis. Chapter 1 gives a very brief account of a bubble chamber, how it is operated, and how the film is used in the reconstruction of events.

The paper tape measuring system is described in some detail in Chapter 2.

The following two chapters describe the on-line system. Chapter 3 is concerned with the design and hardware of the system, and describes in detail the procedure required for the measurement of an event. In Chapter 4 the emphasis is on the computer and how it is programmed. Sections are included on interrupt philosophy, curve fitting and the control of data transfers through the interface.

Finally in Chapter 5 the on-line measuring system is compared with the paper tape system, and several suggestions are given of ways in which the on-line system could be improved.

CONTENTS

	Page
ABSTRACT	i
PREFACE	ii
CONTENTS	v
CHAPTER 1 BUBBLE CHAMBER FILM ANALYSIS	1
1.1 Current Experiments	1
1.2 The CERN 2m Bubble Chamber	2
1.3 The Photographic System	3
1.4 The Bubble Chamber Cycle	4
1.5 A Typical Bubble Chamber Picture	5
1.6 Scanning	7
1.7 Measuring	7
1.8 THRESH and GRIND	9
CHAPTER 2 THE OFF-LINE MEASURING SYSTEM	11
2.1 The Measuring Machines	12
2.1.1 Machines 1 and 2	12
2.1.2 Machine 3	14
2.1.3 Machine 4	16
2.2 The Measurement of an Event	19
2.2.1 The heading information	19
2.2.2 The measurement of a view	21
2.2.3 The completion of an event	23
2.3 REAP	23
2.3.1 Interpretation of different coordinate codes	23

	Page
2.3.2 Detection of tape punching errors	25
2.3.3 REAP output	26
2.4 Problems with the Off-Line System	27
CHAPTER 3 THE ON-LINE MEASURING SYSTEM	31
3.1 General considerations	32
3.2 Choice of computer for the on-line system	34
3.3 The on-line measuring system	37
3.3.1 The computer	38
3.3.2 The measuring machines	40
3.3.3 The OEM Channel	44
3.3.4 The Westinghouse Interface	45
3.3.5 Closed circuit television	50
3.4 The operation of the measuring machines	51
3.4.1 Signing on	52
3.4.2 The start of an event	52
3.4.3 The measurement of a view	56
3.4.4 Signing off	61
3.4.5 Additional features	62
3.5 Preparation of data for further processing	67
CHAPTER 4 PROGRAMMES FOR THE ON-LINE SYSTEM	70
4.1 Introduction	70
4.1.1 A review of the hardware	71
4.1.2 Logic of operation	72
4.1.3 Communication by interrupt	73

	Page
4.1.4 Timing considerations	73
4.2 Interrupts	74
4.2.1 The nature of an interrupt	75
4.2.2 Interrupt levels	76
4.2.3 Interrupt servicing	77
4.2.4 OEM Channel	80
4.2.5 Inhibition of interrupts	80
4.2.6 Operator initiated interrupts	81
4.3 Interrupt Logic	82
4.3.1 Response to an interrupt	83
4.3.2 Selection of data processing routine	84
4.3.3 Data processing routines	85
4.3.4 The OEM channel routine	86
4.3.5 The background programme	91
4.4 Organisation of data in core	94
4.4.1 Common	95
4.4.2 The start of an event	99
4.4.3 The measurement of an event	105
4.4.4 Overlays	107
4.5 Control of data transfers through the interface	109
4.5.1 Details of the OEM channel	109
4.5.2 The Westinghouse interface	114
4.5.3 Programme instructions	119
4.6 Track checking	126

	Page
4.6.1 Circle fitting procedure	128
4.6.2 Apex test	131
4.6.3 Sagitta calculation	132
4.6.4 Precision	133
4.6.5 Coordinate modifications	134
4.6.6 Other curve fitting procedures	136
4.7 Output of data onto disk	138
4.7.1 Organisation of disk storage	138
4.7.2 General considerations	139
4.7.3 Preparation of data for output	141
4.7.4 Continuation addresses	144
4.8 Loading and termination of the on-line system	145
4.8.1 Initialising and loading programme	145
4.8.2 Termination of the on-line programme	147
CHAPTER 5 CONCLUSIONS	148
5.1 Comparisons	148
5.2 Suggested modifications to the on-line system	150
5.2.1 Output of data	150
5.2.2 Inhibition of interrupts	151
5.2.3 Interruption of curve fitting	152
5.2.4 Overlays	153
5.2.5 Fiducial check	155
5.2.6 Curve fitting tolerances	155
5.2.7 Digital output words	156

5.3 Final conclusions	157
APPENDIX A Logic diagram of OEMI routine	159
APPENDIX B Circle fitting procedure	160
APPENDIX C SORT	163
APPENDIX D Console printer messages	167
APPENDIX E List of abbreviations	170
REFERENCES	171
ACKNOWLEDGEMENTS	172

CHAPTER 1

BUBBLE CHAMBER FILM ANALYSIS

Bubble chamber film analysis is one method of studying elementary particle physics. The behaviour of highly energetic charged particles is observed by examining the tracks they leave as they pass through a bubble chamber. This is done by photographing the chamber with three or four cameras simultaneously. A three dimensional reconstruction of the tracks in chamber space is made by measuring the images of the tracks on the film from at least two of the cameras. This reconstruction is a complicated procedure involving the use of iterative methods which can only be carried out on large computers.

Physical results are obtained, not from the detailed observations of individual events, but from the examination of the results obtained by reconstructing many thousands of events. In order to make so many measurements, automatic digitisation is essential.

The operation of the bubble chamber and how the film is used in the reconstruction of the events is outlined in this chapter.

1.1 Current Experiments

The research group in Durham with which the author has been working is currently analysing film from three experiments. These films were exposed to the CERN 2m bubble chamber at the CERN proton synchrotron.

The first of these experiments is an exposure of hydrogen to 12 GeV/c positive pions which was run in July 1966. Work on this



experiment is now nearing completion. About 180,000 pictures were taken, of which 30,000 were allocated to the Durham group for analysis.

The second experiment is the exposure of 12 GeV/c positive pions to deuterium. The films for this experiment were taken in February 1968, and the author was among those who went to CERN at the time to assist in the running of the experiment. In this case Durham's share of the 200,000 pictures was 50,000.

Work has also been started on the scanning and measuring of the film taken during an exposure of deuterium to 4 GeV/c positive pions in August 1970.

1.2 The CERN 2m Bubble Chamber

The films were exposed to the CERN 2m bubble chamber (CERN 2m BC). This consists of an aluminium casting of dimensions 200 x 51 x 60 cm³, with plane glass windows front and back (see Figure 1.1). It is filled with liquid hydrogen, or deuterium depending on the experiment, and is kept under a pressure of about 6 atmospheres at a temperature of about 27°K which is just sufficient to prevent the liquid from boiling. The cycle of operations will be described in Section 1.4.

The chamber is surrounded by a large electromagnet which produces a nearly uniform magnetic field of about 17 Kgauss over the whole chamber. The direction of the field is normal to the windows of the chamber.

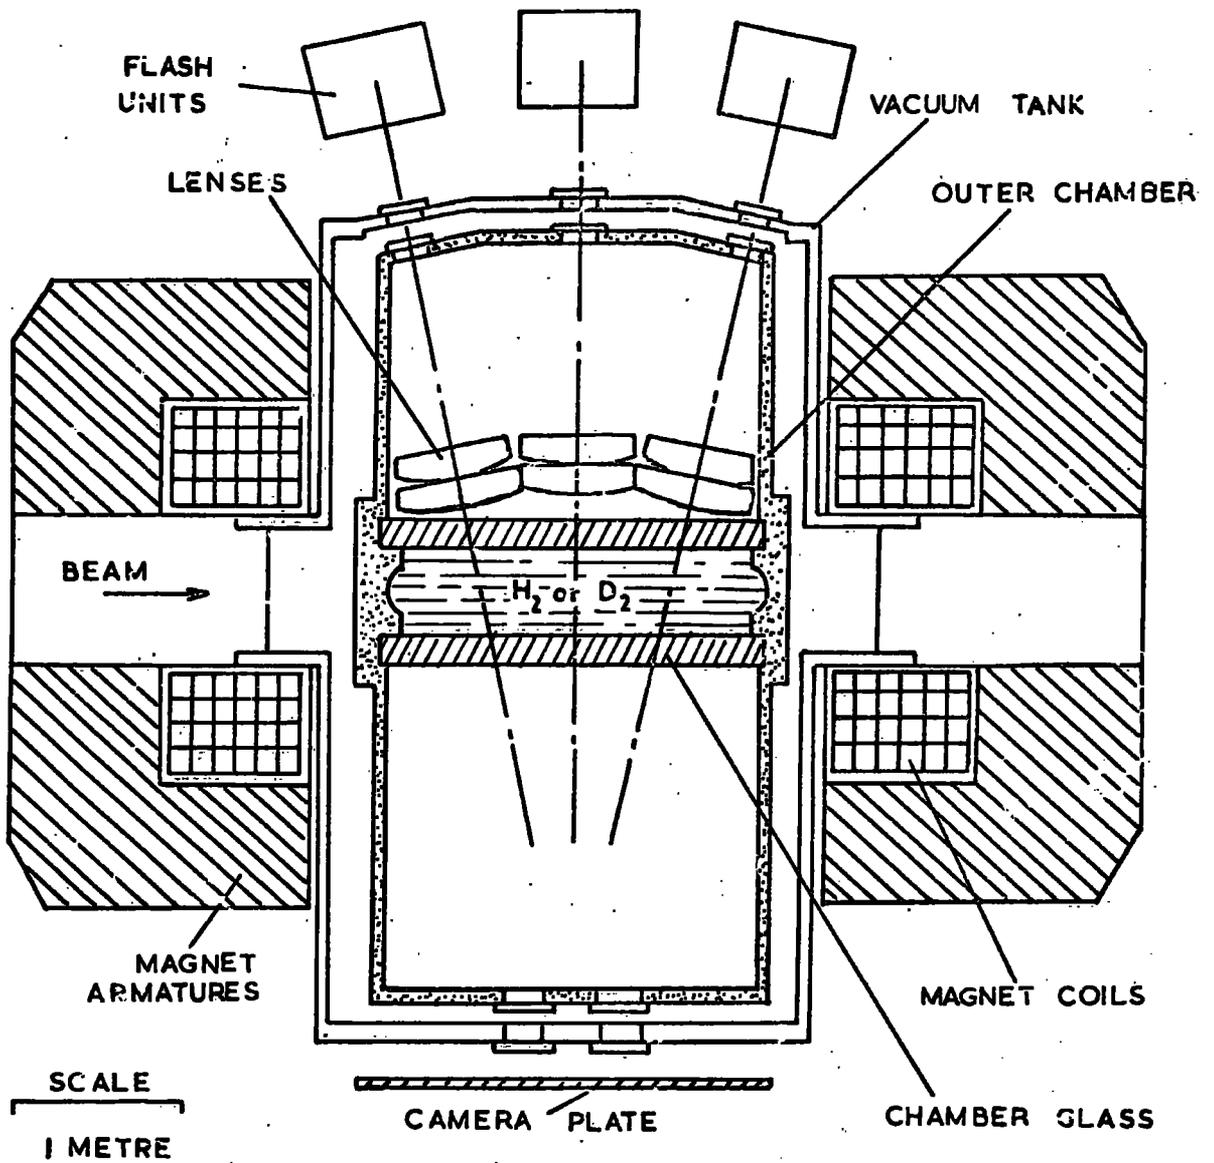


FIGURE 1.1 - CERN 2 METRE BUBBLE CHAMBER -

Four cameras are arranged so that the whole chamber can be photographed stereoscopically. Fiducial marks are engraved on the inside of both the front and the back windows of the chamber. These provide a coordinate system which is independent of the photographic system.

1.3 The Photographic System

The CERN 2m BC uses the through illumination type of optical system. Three flash tubes are used with monochromatic filters in the range 4200 - 4600 Å.

The four cameras are mounted in a diamond shaped array with their axes perpendicular to the chamber windows. Their relative positions and focal lengths are shown in Figure 1.2 and Table 1.1.

TABLE 1.1

Focal lengths and relative positions of the cameras

CAMERA	FOCAL LENGTH	POSITION
1 TOP	181.5 ± 0.7 mm	x = 410 ± 0.1 mm y = 600 ± 0.1 mm
2 ENTRY	182.5 ± 0.7 mm	x = 110 ± 0.1 mm y = 300 ± 0.1 mm
3 EXIT	182.1 ± 0.7 mm	x = 710 ± 0.1 mm y = 300 ± 0.1 mm
4 BOTTOM	183.3 ± 0.7 mm	x = 410 ± 0.1 mm y = 0 ± 0.1 mm

For each camera the distance between the first principal plane (object side) and the inner face of the front chamber window (the reference plane)

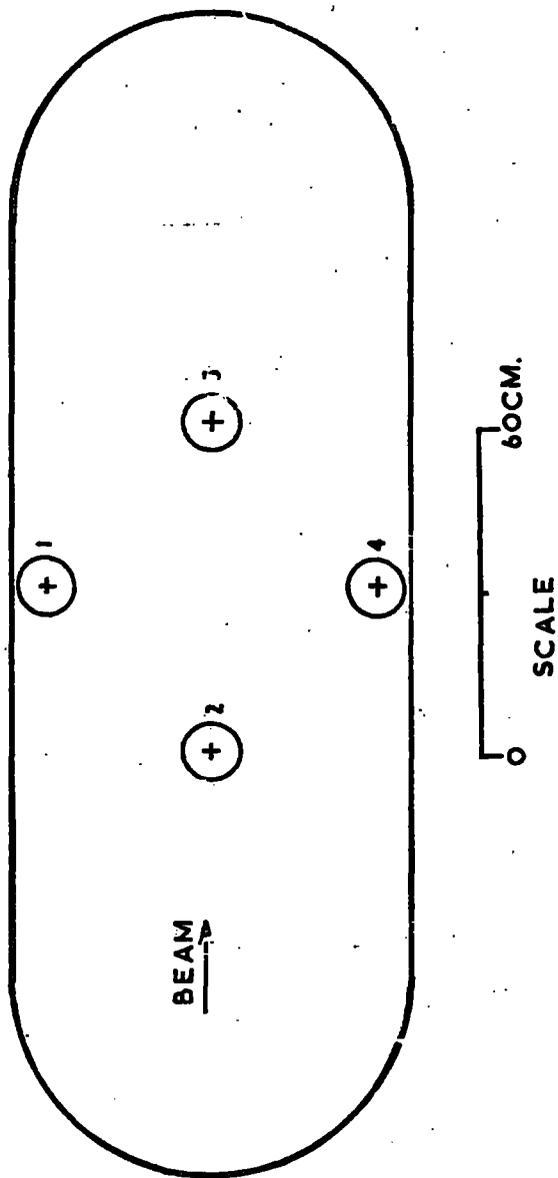


FIGURE 1.2 —CAMERA POSITIONS OF THE CERN METRE CHAMBER —

is 2430 ± 5 mm. The camera plate and the film planes are parallel to the reference plane to within ± 0.5 milliradian. Each camera has an aperture of $f/32$ which gives a depth of field which covers the whole chamber from front to back.

Each camera views through a set of three small windows which are arranged in such a way that the deviations produced by the different windows partly compensate each other. The resulting deviation is no more than 0.15 ± 0.1 milliradian.

1.4 The Bubble Chamber Cycle

The operation of the bubble chamber is designed so that charged particles passing through the chamber leave a trail of bubbles which can be photographed.

Initially the pressure on the bubble chamber is just sufficient to prevent the liquid from boiling, and when this is rapidly reduced from about 6.5 to 2.8 atmospheres, the liquid hydrogen (or deuterium) becomes superheated and spontaneous boiling occurs at points of high charge density. When the chamber is in this sensitive condition a beam of highly energetic positively charged pions is admitted. These produce ionisation which causes spontaneous boiling to occur along the path of the particle. This appears as a trail of bubbles which are allowed to grow until they are about 200μ in diameter before the chamber is illuminated, and photographed. Before spontaneous boiling occurs throughout the chamber, the pressure is increased again, the

bubbles disappear, and the chamber is allowed to settle before the next expansion. The complete cycle is repeated every two seconds.

The bubble chamber is usually operated for about a week at a time, resulting in the production of 100,000 to 200,000 pictures. The films are then shared between the different laboratories taking part in the experiment.

More recent experiments use the technique of double pulsing which results in the production of 350,000 pictures per week.

1.5 A Typical Bubble Chamber Picture

Figure 1.3 shows two views of a typical bubble chamber picture taken from an exposure of deuterium to 4 GeV/c positive pions. Slight differences due to the separation of the cameras 1 and 2 which were used can easily be distinguished. In particular the relative positions of the fiducial marks may be seen to depend on the view.

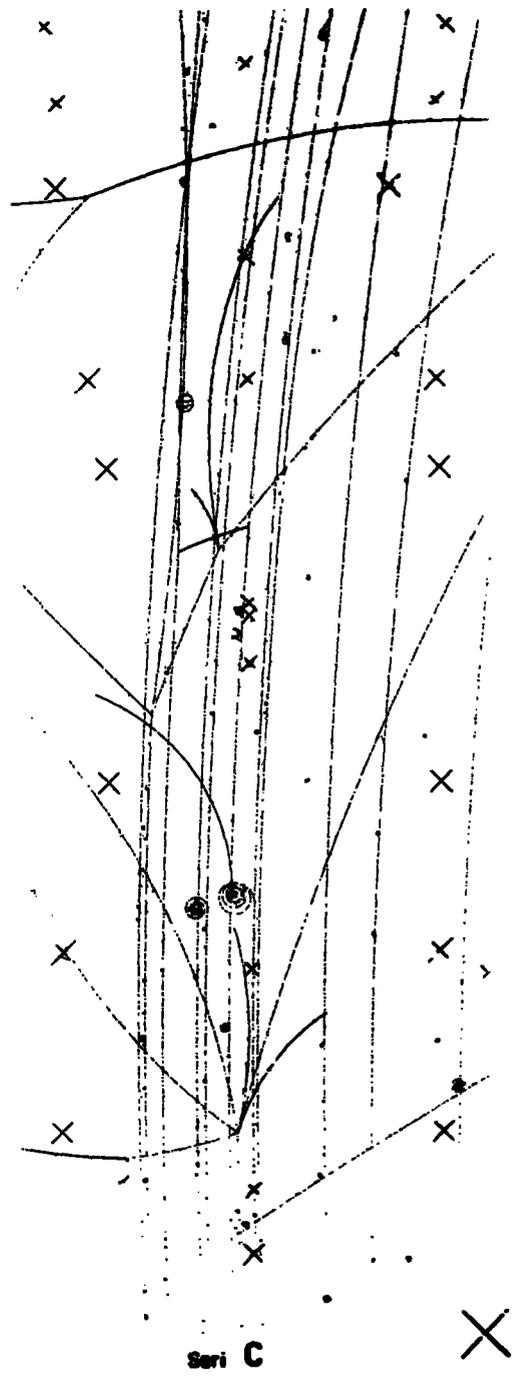
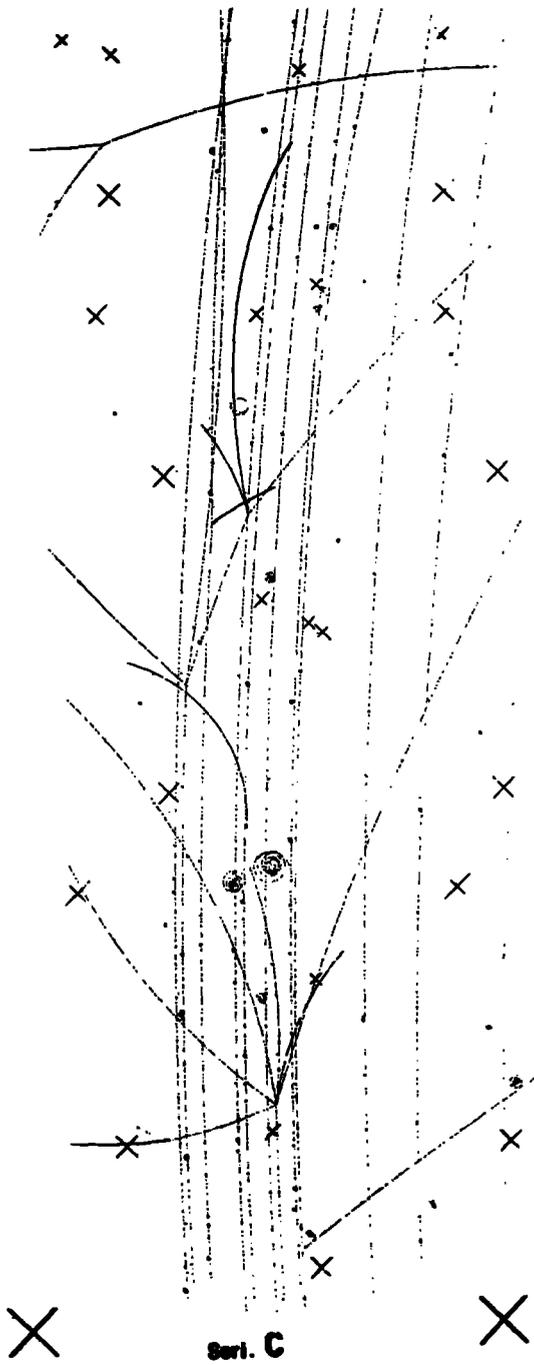
The tracks of all the particles are curved due to the magnetic field. A particle which traverses the chamber with no significant loss of energy will leave a track in the form of a helix. Since the axes of the cameras are parallel to the magnetic field such a track will appear in each view as a close approximation to an arc of a circle. When the loss of energy is significant the track is in the form of a spiral, which in projection, only approximates to an arc of a circle over very short lengths, because of the increasing curvature.

In the two views in Figure 1.3 the particles are shown entering

Figure 1.3 - A typical bubble chamber picture

VIEW 1

VIEW 2



the chamber at the foot of the page. All the particles in the beam have the same incident momentum of 4 GeV/c and so leave distinctive parallel tracks in the chamber.

Six secondary tracks can be seen leaving the apex of the first event in this picture. These were produced by four positively and two negatively charged particles as they left the interaction. On one of the positive tracks a secondary two prong event can be seen.

The second primary event appears as a three prong interaction half way up the picture. Close observation will reveal another very short dark track, which was formed by a proton which received so little energy that it was brought to rest almost immediately. This event would then be classed as a four prong event with a very short stopping proton. A three prong secondary event appears on one of the positive tracks.

Another stopping track which may be seen on the next event, was produced by a proton which lost all its energy before leaving the bubble chamber. By measuring the stopping point of the track an accurate evaluation of the momentum of the particle can be made from the range energy relation. For a track which leaves the chamber, the evaluation of the momentum has to be based on a measurement of the curvature of the track.

A bubble chamber picture shows only the tracks of charged particles as uncharged particles do not produce ionisation and so leave no tracks. However, their presence can sometimes be detected in the bubble chamber picture. If an uncharged particle decays within the chamber into two charged particles, the tracks of these two particles will be seen. They

will appear in the form of a V which points towards the interaction in which the uncharged particle was formed.

The picture may also reveal distinctive characteristics of some charged particles. For instance, certain "strange" particles which are charged may decay within the chamber into two new particles of which only one is charged. As the uncharged particle leaves no track it will appear as if there was a "kink" on the track of the original particle.

1.6 Scanning

Before starting to measure any film it is scanned for events. Three views are examined, and all events found are listed, a record being made of their location, size and any special features such as stopping tracks or close secondaries. The size of an event is determined by the number of secondary tracks leaving the apex of the event.

A second scan is carried out and the two lists compared and checked against the film. If the scanning accuracy is high a single scan may be considered sufficient.

From the scan sheets measuring lists are prepared.

1.7 Measuring

For any particular track the measurements from any two views are sufficient for the three dimensional reconstruction in chamber space, provided the track does not lie along the axis of either camera. To avoid such a bad view of any track, three views are measured, thus

ensuring at least two good views of each track. The slight improvement in the results obtained by measuring four views does not justify the extra work required.

The coordinate system of the measurements is provided by measuring at least four of the fiducial marks. Of these it is desirable that at least two should be on the front window and at least two on the back window of the chamber. It is then possible to check that each view is correctly identified, since the relative positions of the fiducial marks are peculiar to each view. The measurements of the fiducial separations are also used to determine the extent of lateral and longitudinal stretching of the film which may vary not only from one roll of film to another, but also along each roll.

For each track usually nine points are measured in each view, although short tracks may be accepted with as few as five points. For very short tracks only the end point is measured. Thus a simple two prong event requires the measurement of nearly 100 coordinate pairs, and a four prong event requires about 150 coordinate measurements. These measurements are listed in Table 1.2.

Besides the recording of the coordinate pairs for the different points measured, all the measurements have to be labelled so that the related measurements from the different views can later be identified and combined in the reconstruction process. The event itself is identified by a set of eight digits made up of the film number, the frame number and the event number. Each view is identified and each point or track is given a label.

TABLE 1.2

Number of coordinate measurements required in each view for
a 2-prong and for a 4-prong event

MEASUREMENT	2-PRONG EVENT	4-PRONG EVENT
Fiducial marks	4	4
Apex	1	1
Beam track	9	9
Secondary tracks	18	36
View Totals	32	50

The data are then passed on to the programmes THRESH and GRIND for the reconstruction of the event.

1.8 THRESH and GRIND

The geometrical reconstruction of the events is carried out by the programme THRESH which is run on the Northumbrian Universities Multiple Access Computer. NUMAC consists essentially of an IBM 360/67 computing system sited in Newcastle, with an IBM 1130 used as a terminal computer in Durham.

THRESH first of all checks the measurements of each view in turn. The fiducial measurements are examined to see if they lie within the permitted tolerance of their expected positions. A two dimensional circle

fit is made to the measurements of the nine or so points along each track. These best fit circles are then used to reconstruct the tracks in three dimensional chamber space. The two best views for each track are used for the reconstruction, and the third view, if present, is used for checking.

Events which are successfully reconstructed in THRESH are passed on to GRIND which calculates the probabilities that the different hypotheses are good descriptions of each event.

Any events which are rejected by THRESH or GRIND are remeasured and the new measurements processed by THRESH and GRIND in the usual way.

CHAPTER 2

THE OFF-LINE MEASURING SYSTEM

The measurement of several thousands of bubble chamber events requires the use of some kind of digitising machines which determine the coordinates of points along tracks on the film both accurately and rapidly. A precision of a few microns on the film is required. The measurements must be made available in a form that can be input to a computer so that a three dimensional reconstruction of each event can be made. Some way of identifying the measurements with particular tracks and events is also required.

The digitising machines used by the Durham group are described in this chapter. Before the on-line measuring system was built, the measurements and their associated labels were punched onto paper tape. The data were later edited and sorted by a programme REAP, and were output onto cards in a form suitable for input to the geometry programme THRESH.

Initially the editing was carried out on an Elliott 803 computer in Durham and its paper tape output was processed by THRESH and GRIND in the IBM 7044 computer then being used by the Film Analysis Group at Glasgow University. When this chain of processing was being converted to an IBM 1130 computer in Durham and the IBM 360/67 at Newcastle, a new version of the editing programme REAP was written by the author for the IBM 1130 computer.

This chapter describes the measuring machines and how they were operated with the tape punching equipment. Some of the advantages and disadvantages of this system are discussed.

2.1 The measuring machines

Four measuring machines were used in the off-line measuring system. Three of these were digitised in the film plane, and the fourth used an image plane coordinatograph.

The coordinate measurements from each of the four machines were punched onto eight hole paper tape. Heading information and labels which were required to identify the measurements were entered using input writers, and were also punched onto the tape.

The tape from all of the machines was punched in even parity. This provided a quick and simple check on the tape punching machines as the operators could notice at a glance if there were any odd parity characters on the tape.

2.1.1 Machines 1 and 2

On the first two machines, the films of the three views being measured are mounted so that each in turn can be projected onto a screen with a magnification of about $\times 15$. The films are clamped to a stage which can be driven in the X direction under the control of the operator. This forms the X substage. The lamp and condenser mounted above and below the film respectively form the Y substage, and can be driven in the Y direction, independently of the X movement, also under operator control.

A reference mark, about the size of a bubble, is marked on the screen. Measurements are made by driving the two substages so as to bring the image of the film into coincidence with this mark. The mark

lies close to the optic axis of the projection system, so that distortions due to projection are minimised.

Since the film itself is moved to make the measurements in the X direction, these measurements are of the film movement. This is illustrated in Figure 2.1. In the Y direction however, it is the movement of the lens rather than that of the film which is measured, and this is shorter than the apparent film movement by about 8%. The Y movement is illustrated in Figure 2.2. The geometry programme THRESH handles the different magnifications in the X and Y directions by calculating these magnifications from the measured positions of the fiducials.

Each of these machines uses a Moiré fringe optical digitisation system. This consists essentially of two gratings, one mounted on the moving part of the substage, the other fixed to the main part of the machine, which produce an interference pattern whose fringes are detected by photoelectric cells. The precision of the machines is limited by the fringe size which corresponds to about 2.5μ on the film.

On the film the bubbles are $12-14\mu$ in diameter. With care it is possible to set on a track to about half a bubble diameter, which corresponds to about 2 to 3 fringes. Under normal working conditions it is found that setting accuracy is about 4 fringes, which corresponds to 10μ on the film.

The photoelectric cell outputs are fed into two 17 bit bi-directional binary counters, one for the X coordinates and one for the Y. These are displayed to allow the operator to check the counting easily.

The measurements were punched onto paper tape, and each pair of coordinates appeared as six characters preceded by a coordinate starting

TO BRING Q, THE IMAGE OF q, INTO COINCIDENCE WITH REFERENCE MARK R, THE FILM IS MOVED A DISTANCE x .

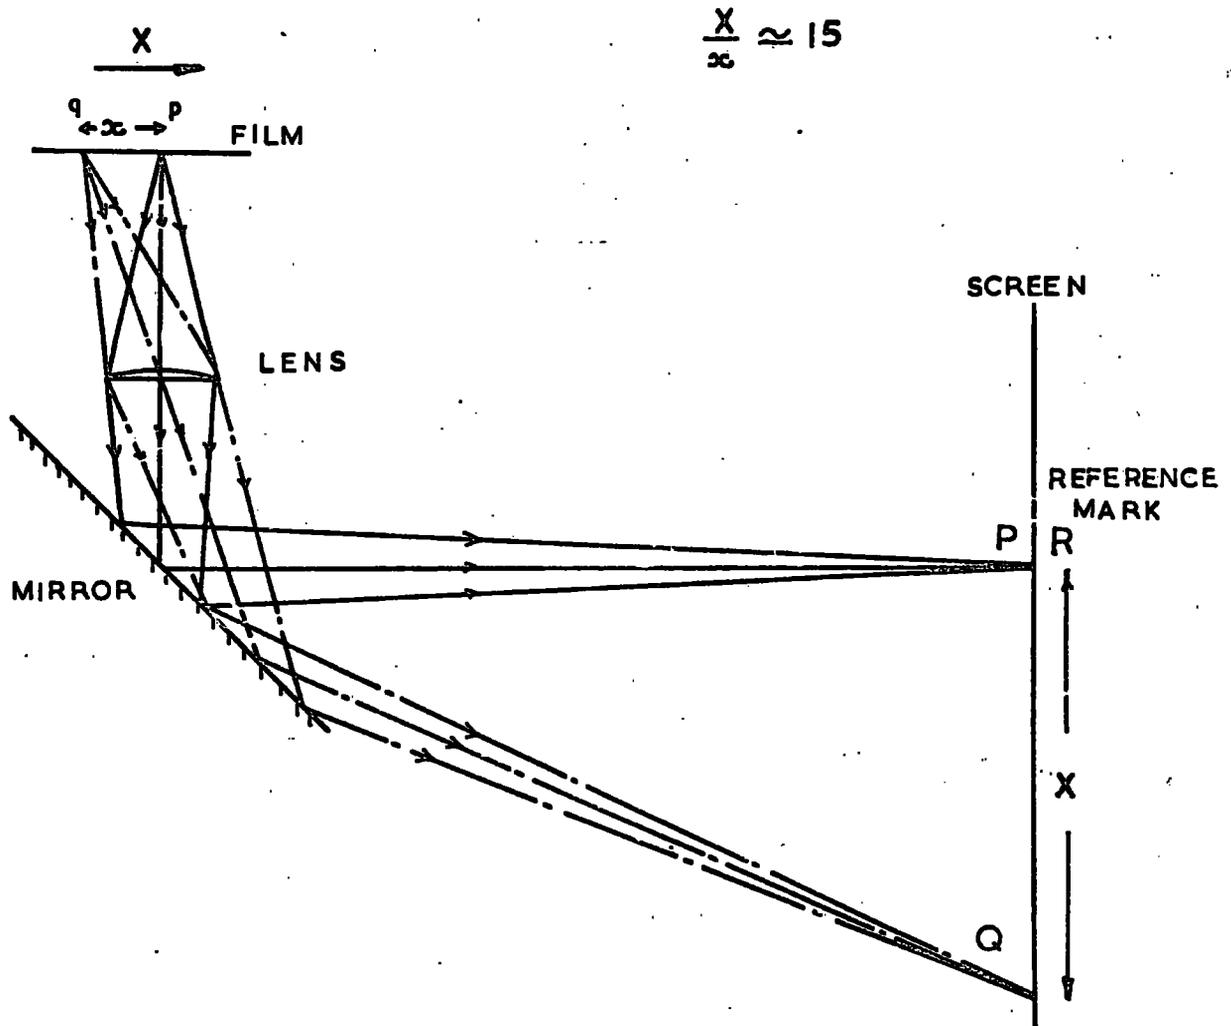


FIGURE 2.1-OPTICAL SYSTEM OF MACHINES 1 & 2, X MOVEMENT -

TO BRING Q, THE IMAGE OF q, INTO COINCIDENCE WITH REFERENCE MARK R, THE LENS IS MOVED A DISTANCE y.

$$\frac{y}{y'} = \frac{a+b}{b} \approx 1.08$$

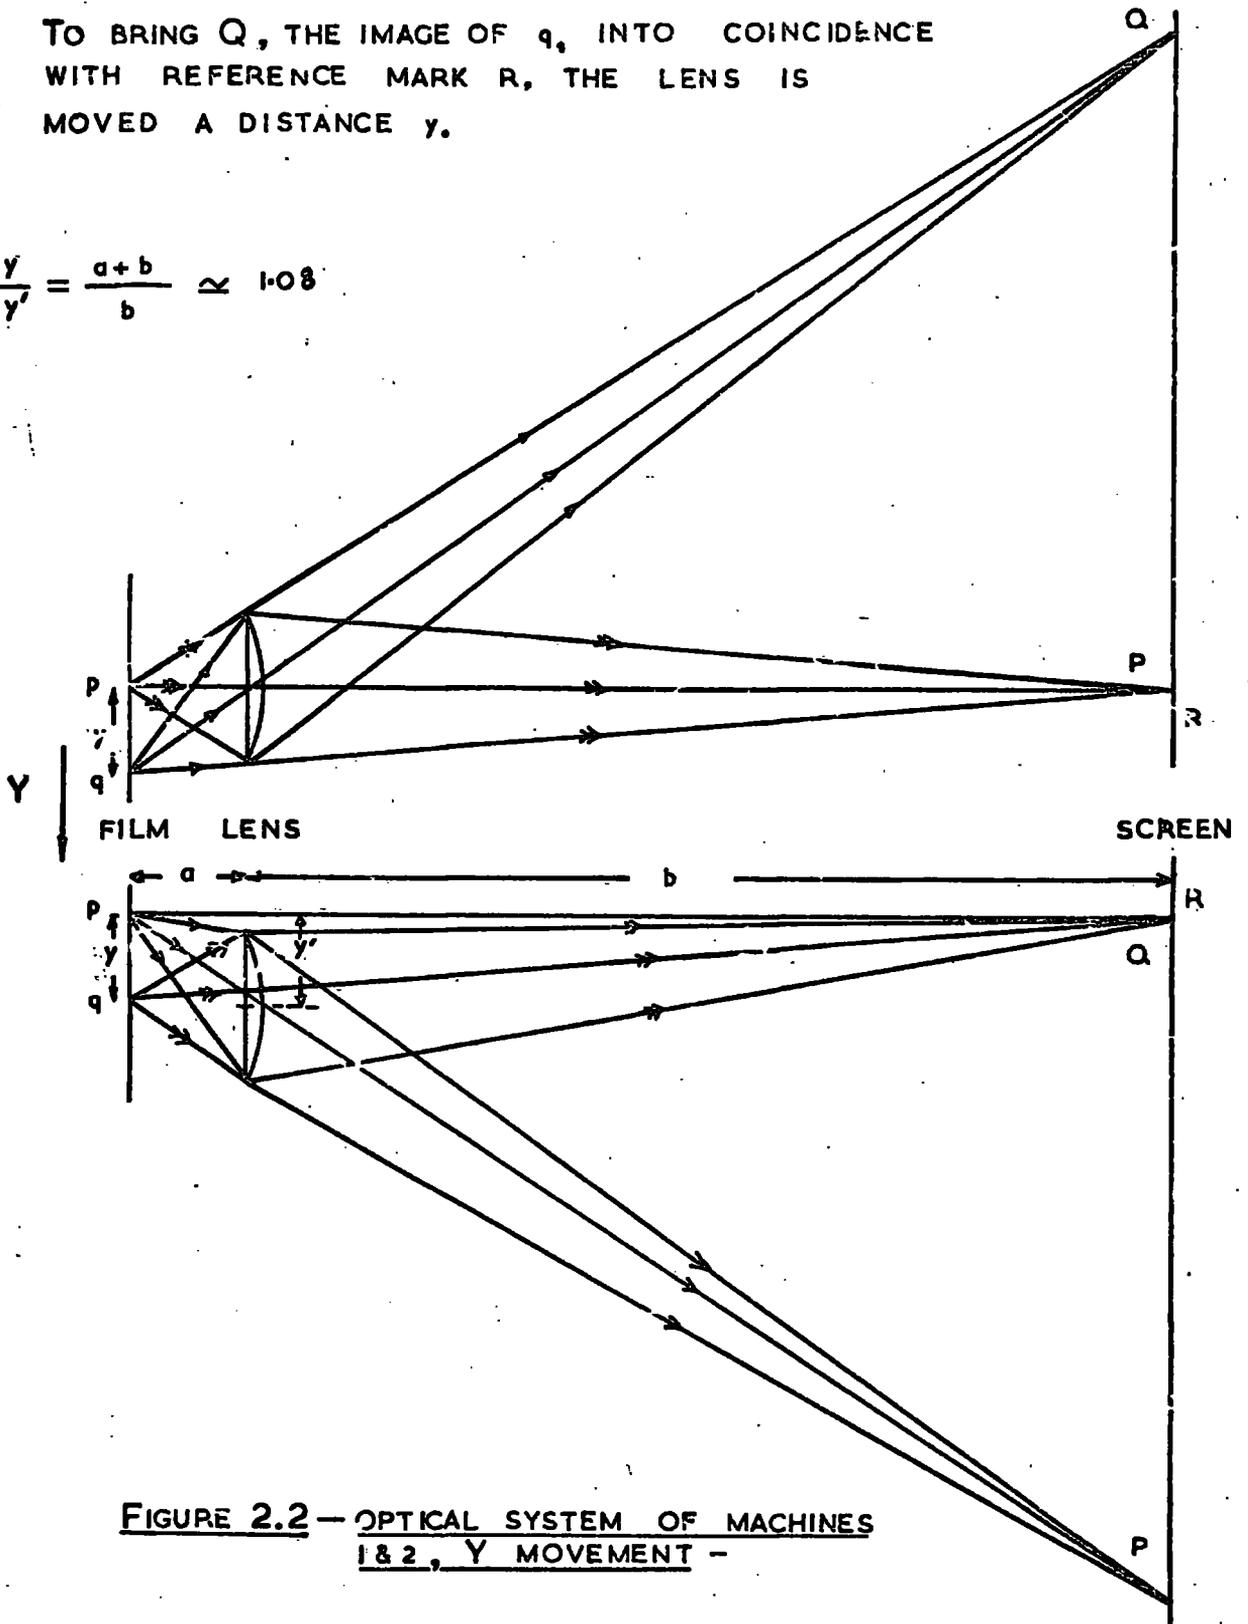


FIGURE 2.2 — OPTICAL SYSTEM OF MACHINES 1 & 2, Y MOVEMENT —

character. At the same time a "/" character was printed on the input writer to enable the operator to keep account of the number of points measured. Figure 2.3 shows the relationship between the binary value in the coordinate counters and the six characters that were punched for a typical pair of coordinates.

2.1.2 Machine 3

Machine 3 is like the first two machines in that it is also digitised in the film plane. In this case however, the stage carrying the three films of the three views being measured can be moved in both the X and the Y directions, so that the measurements made in each direction are of the film movement itself.

Each view in turn is projected onto a screen. Measurements are made by driving the stage, using the X and Y controls, in order to bring the image of the film into coincidence with the reference mark on the screen.

Instead of the Moiré fringe system of the other machines, this machine was fitted with a mechanical digitization system. Two Hilger and Watts mechanical digitisers were used to measure the rotation of the machine drive shafts of the motors which moved the stage. Each consisted of a drum mounted on the shaft, coded into conducting and non-conducting segments, which were sampled by brush contacts. A reflected decimal code was used which had the advantage that in moving from one digit to another

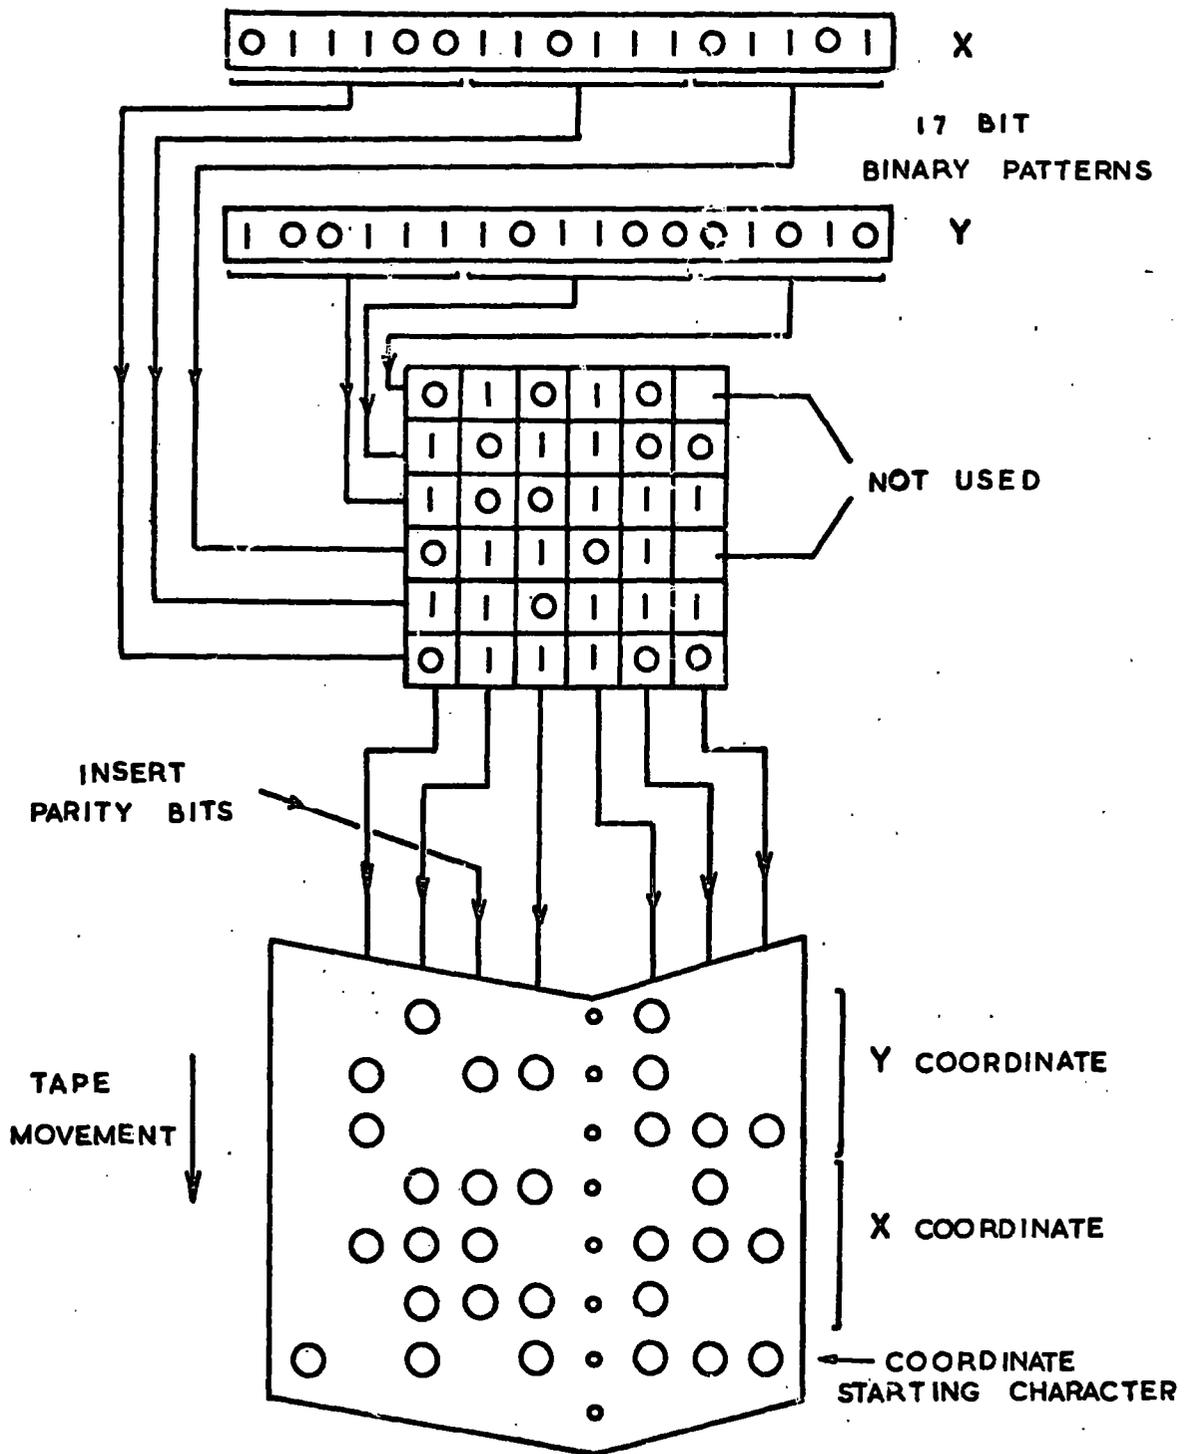


FIGURE 2.3 - MACHINES 1 & 2 COORDINATE CHARACTERS -

only one segment changed, and so errors due to the brush contacts being imperfectly aligned were eliminated.

The precision of the machine was limited by the width of the segments, which corresponded to about 2μ movement of the film. The magnification of the projection system is about the same as for the other film plane digitised machines. Under normal working conditions points are measured to within about 10μ on the film.

As for the other machines, the output from the digitisers was punched onto paper tape. In this case, 11 characters were punched for each coordinate pair. These consisted of the starting character followed by two five digit decimal numbers. A typical example of the 11 characters that were punched for a single pair of coordinates is illustrated in Figure 2.4. Table 2.1 shows the reflected decimal coding which was used for this machine.

Table 2.1

Reflected decimal code

VALUE	NORMAL	REFLECTED
0	0 0 0 0	1 0 0 0
1	0 0 0 1	1 0 0 1
2	0 0 1 1	1 0 1 1
3	0 0 1 0	1 0 1 0
4	0 1 1 0	1 1 1 0
5	1 1 1 0	0 1 1 0
6	1 0 1 0	0 0 1 0
7	1 0 1 1	0 0 1 1
8	1 0 0 1	0 0 0 1
9	1 0 0 0	0 0 0 0

EACH OF THE LAST THREE DIGITS FOR EACH COORDINATE IS PUNCHED IN NORMAL OR REFLECTED CODING DEPENDING UPON WHETHER THE PRECEDING DIGIT REPRESENTED AN EVEN OR AN ODD NUMBER RESPECTIVELY. THE FIRST TWO DIGITS ARE ALWAYS NORMAL.

e.g. X = 25861

2	○ ○	NORMAL
5	○	NORMAL
8	○ ○ ○	REFLECT.
6	○ ○	NORMAL
1	○ ○ ○	NORMAL

Y = 12754

1	○ ○ ○	NORMAL
2	○ ○	NORMAL
7	○	NORMAL
5	○ ○	REFLECT.
4	○	REFLECT.

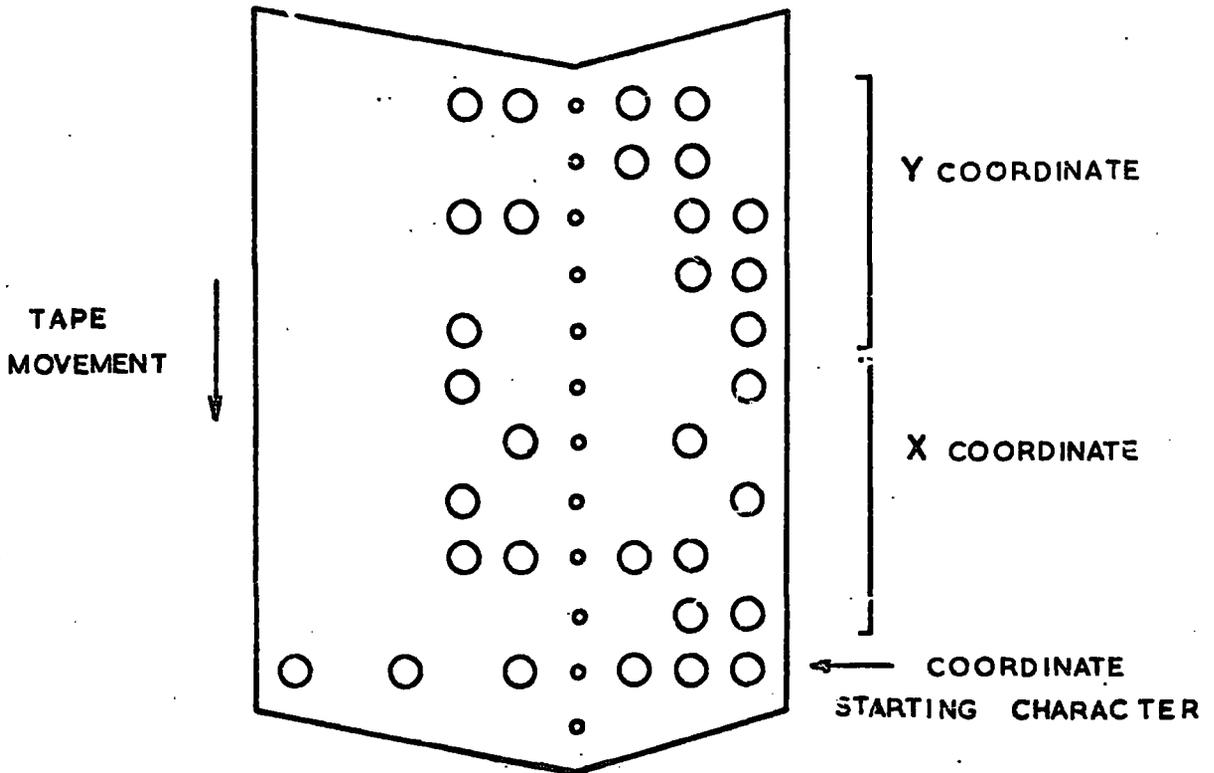


FIGURE 2.4 — MACHINE 3 COORDINATE CHARACTERS —

As usual the "/" character was printed on the input writer for each coordinate pair that was punched onto paper tape.

2.1.3 Machine 4

The fourth machine is digitised in the image plane rather than the film plane. Three projectors are mounted so that each in turn can project its view onto the measuring table top with a magnification of about $\times 10$. Instead of moving the film or the lens system to obtain coincidence of image and reference mark, the mark itself, which is mounted on a coordinatograph system, is moved about over the table top. It is this movement of the mark which is digitised.

The coordinatograph has a Moiré fringe digitisation system whose counters are displayed as two 6 digit octal numbers. The fringe width, which limits the precision of the measuring machine, corresponds to 25μ on the table top, which is equivalent to about 2.5μ on the film.

In this case it is not possible to confine measurements to a region close to the optic axis of the measuring system. In particular, the fiducial marks which always have to be measured, are widely separated, and on average will be at least 7 degrees from the optic axis. Lens distortions at such angles are not insignificant, and as the reconstruction of events is very sensitive to the measured positions of the fiducials it is important to make corrections for these distortions.

There are also obliquity distortions arising from the three projectors producing an image on the same table top. These are reduced as far as possible by arranging the middle of the three projectors so that its optic axis is approximately perpendicular to the table top. The distortions on the other two views are greater. Again the fiducial measurements are badly affected by these distortions.

The corrections in THRESH allow for uniform stretching across the width and along the length of the film, but do not allow for the non-uniform distortions produced by this projection system. It is therefore necessary to make certain modifications to the coordinate measurements before the data are input to THRESH.

The modifications made are based on a comparison between coordinate measurements obtained from this machine and those obtained from machine 3. Machine 3 is a true film plane digitised machine, and measurements made on that machine are free from distortions due to the projection system. A set of 20 fiducial marks was measured on both machines 3 and 4. A linear regression programme/^{was} used to calculate the coefficients of an expression that would transform the coordinate measurements from machine 4 into the corresponding measurements for machine 3. This was done separately for each of the three views and the conversion expressions which were calculated in this way are shown in Table 2.2.

Care has to be taken to ensure that the film is always projected onto the same part of the table top as the distortions depend upon the absolute position of the image on the table. This is achieved by insisting that fiducial 33 is always projected onto the same position on the table top.

Table 2.2

Expressions to convert machine 4 coordinate measurements
into machine 3 equivalents

VIEW	CONVERSION EXPRESSION
1	$X_3 = -4773.1 + 1.143 X_4 - 0.0325 Y_4 + 1.0 \cdot 10^{-6} X_4 Y_4$ $Y_3 = -2988.3 - 0.0245 X_4 + 1.077 Y_4 + 1.2 \cdot 10^{-6} Y_4^2$
2	$X_3 = -5660.8 + 1.167 X_4 - 0.0137 Y_4 + 0.6 \cdot 10^{-6} X_4 Y_4$ $Y_3 = -4642.7 - 0.0234 X_4 + 1.147 Y_4 + 0.6 \cdot 10^{-6} Y_4^2$
3	$X_3 = -7425.1 + 1.2157 X_4 + 0.0471 Y_4 - 1.1 \cdot 10^{-6} X_4 Y_4$ $Y_3 = -6098.1 - 0.025 X_4 + 1.2484 Y_4 - 1.1 \cdot 10^{-6} Y_4^2$

Note: X_1, Y_1 measured in fringes

By modifying the coordinate measurements from machine 4 in this way, point reconstruction errors in THRESH were reduced by a factor of four. An earlier attempt to correct for the obliquity distortions based on calculations of the geometry of the projection system was found to be unsatisfactory.

The output from machine 4 was punched onto paper tape, but in this case the usual coordinate starting character was followed by two 6 digit numbers, corresponding to the 6 digit octal values of the coordinate

measurements. A typical example of a pair of coordinates is illustrated in Figure 2.5.

2.2 The measurement of an event

In order to measure an event the coordinates of the apex, four fiducials, and about 9 points on each track were punched onto paper tape. This was repeated for each of the three views. These coordinates were preceded by labels identifying the point or track to which they referred. Certain heading information was also required to identify the event, the operator, the date and the experiment.

The following paragraphs describe in detail the procedure that was required to measure each event.

2.2.1 The heading information

Each event was started by the "<" character being entered on the input writer. This punched one character on paper tape which was peculiar to the machine being used, and which enabled the editing and sorting programme REAP to identify the machine.

After the opening bracket, an eight digit number was punched, to identify the film, the frame and the event. Then the operator number was entered followed by the topology, a five digit number describing the type of event about to be measured. The operator then entered the date, the machine number, and the measurement number.

$X = 124760_8$ $Y = 057213_8$

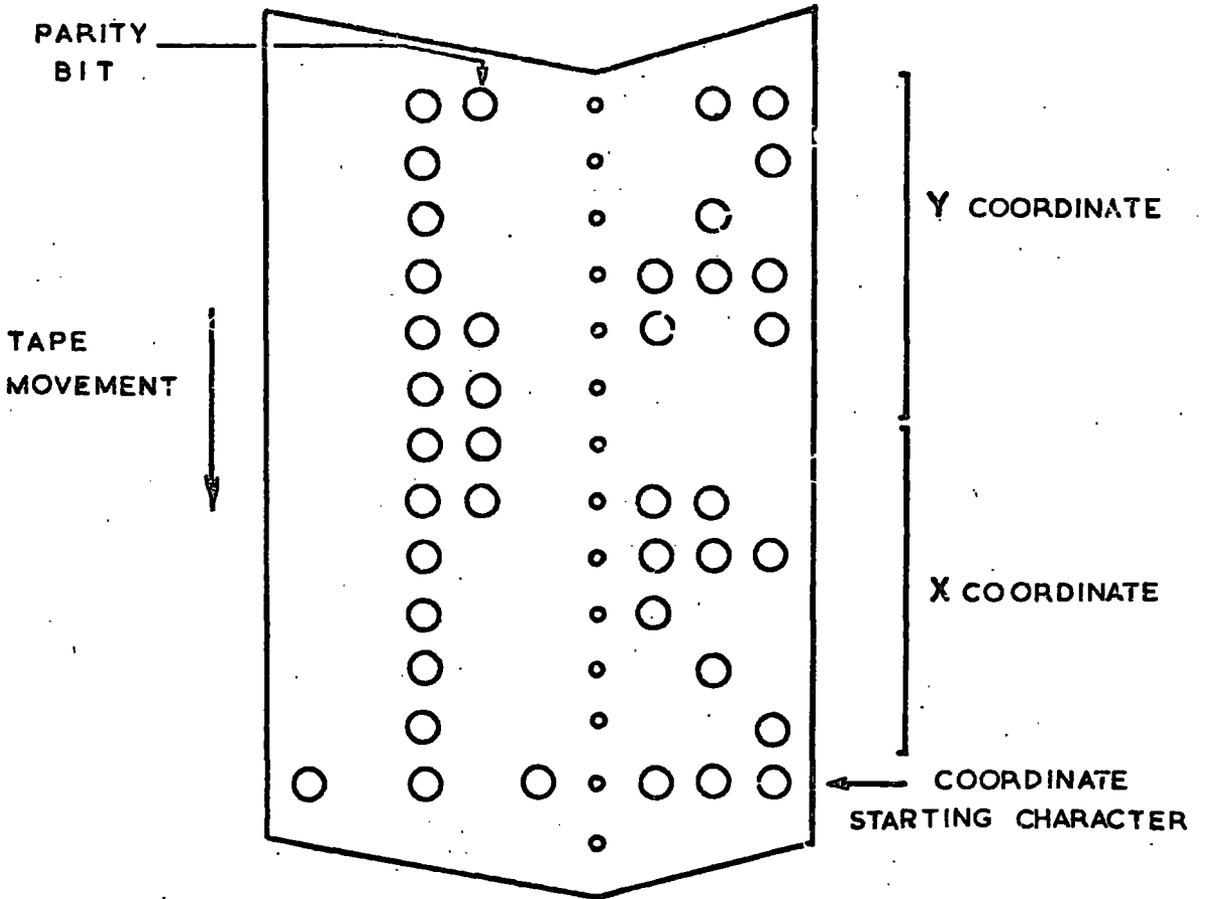


FIGURE 2.5 - MACHINE 4 COORDINATE CHARACTERS -

A full set of heading information was required for each event. Sometimes, however, only the event number changed from one event to another, in which case only the event number needed to be entered. The rest of the heading information was carried over from the previous event by REAP. Similarly, if the date, machine number and measurement number were to be the same, these could be omitted, requiring only the event number, operator number and topology to be punched.

Yet another option was included in this system, which was used when two or more events were to be measured on the same frame. In this case the operator could use the letter "C" instead of the usual opening bracket "<" to mark the start of the second or subsequent event. This meant that REAP would preserve the measurements of the fiducial marks from the first event and use them for the second event, thus relieving the operator of the necessity of measuring the fiducial marks again. Care was required over the setting of the coordinate counters when this option was used. Usually the counters were set at the beginning of the measurement of each view when the reference mark was in coincidence with the apex of the event. The measurements of the fiducials from the first event on the frame could only be valid for the second event if the counters were still counting relative to ^{the} apex of the first event. This required that, during the measurement of the second event, the operator set the counters at the beginning of each view not on the apex of the event being measured, but on the apex of the first event to have been measured in that frame.

Some examples of the use of these different options can be seen in Figure 2.6 which is a sample listing from the machine 1 input writer.

Figure 2.6 Sample listing from machine 1 input writer

```
<38100622 13 02010 10/12/68 1 1=  
+1aa/a9////////11/44/33/22/a1////////a2////////a3////////  
+2aa/a9////////11/44/33/22/a1////////a2////////a3////////  
+3aa/a9////////11/44/33/22/a1////////a2////////a3////////>  
<38100471 13 03010=  
+1aa/a9////////11/44/33/22/xx/a2////////a1////////a3////////  
+2aa/a9////////11/44/33/22/xx/a2////////a1////////a3////////  
+3aa/a9////////11/44/33/22/xx/a2////////a1////////a3////////>  
<38100451=  
+1aa/a9////////11/44/33/22/ap/////pp/a2////////a1////////a3////////  
+2aa/a9////////11/44/33/22/ap/////pp/a2////////a2a1////////a3////////  
+3aa/a9////////11/44/33/22/ap/////pp/a2////////a1////////a3////////>  
c38100452=  
+1aa/a9////////a1////////a3////////a2////////a5////////  
+2aa/a9////////a1////////a3////////a2////////a5////////  
+3aa/a9////////a1////////a3////////a2////////a5////////>  
<38100231=  
+1aa/a9////////11/44/33/22/a1////////a2//a1a3////////*/a2////////pp/ap/////   
+2aa/a9////////11/44/33/22/a1////////a3////////a2////////pp/ap/////   
+3aa/a9////////11/44/33/22/a1////////a3////////a2////////pp/ap/////>  
<38100201 31 03010 11/12/68 1 1=  
+1aa/a9////////11/44/a1////////a2////////a3////////22/33/a5////////  
+2aa/a9////////11/44/a1////////a2////////a3////////a3////////a5////////  
22/33/  
+3aa/a9////////11/44/a1////////a2////////22/33/a3////////a5////////>  
c38100202=  
+1aa/a9////////a1////////a2////////a3////////a5////////  
+2aa/a9////////a1////////a2////////a3////////a5////////  
+3aa/a9////////a1////////a2////////a3////////a5////////>  
c38100204=  
++1aa/a9////////a1/*////////a2////////a3////////xx/  
+2aa/a9////////a1////////a3////////a2////////xx/  
+3aa/a9////////a1////////a3////////a2////////xx/>  
<38100071=  
+1aa/a9////////11/44/a1////////a2////////a3////////22/33/ap/////pp/  
+2aa/a9////////11/44/a1////////a2////////a3////////22/33/ap/////pp/  
+3aa/a9////////11/44/a1////////a2////////a3////////22/33/ap/////pp/>  
c38100073 31 02010=  
<38100043 31 02010=  
+1aa/a9////////11/44/a1////////a2////////a3////////22/33/  
+2aa/a9////////11/44/a1////////a3////////a2////////22/33/  
+3aa/a9////////11/44/a1////////a3////////a2////////22/33/>
```

The different numbers making up the heading information were separated by space characters. The end of the heading information was indicated by the "=" character, or by the "+" which ^{was} used as the view starting character.

2.2.2 The measurement of a view

Throughout the measurement of each event a list of the measurements made was printed by the input writer. An example of such a list is shown in Figure 2.6. The "/" character was printed whenever a pair of coordinates was punched onto paper tape. All the other characters were entered on the input writer itself.

The start of each view was denoted by the "+" character followed by the view number. The counters on the digitisers were set by the operator when the apex of the event was in coincidence with the reference mark. This setting value was in the middle of the range of the counters, e.g. 2^{16} for the 17 bit binary counters. This ensured that all the coordinates entered would be positive, so avoiding the need for a sign bit. No setting of the machine 3 mechanical digitisers was required as the range of the 5 digit decimal numbers was sufficient to cover all three views.

The operator then proceeded to measure the event. A label was entered before each measurement was made to identify the point or track to which it referred. For example "AA" would have been entered before the coordinates of the apex of the event were punched, or "A9" would have been followed by 8-10 pairs of coordinates for points lying along the beam track. These measurements could be made in any order.

The setting value was checked at frequent intervals throughout the measurement of each view, by returning to the apex and checking its characteristic coordinate of $(2^{16}, 2^{16})$. If at any time this was found to be in error, as could easily be seen from the display of the counters, the counters were reset, and the sequence of measurements made since the previous acceptable check was repeated.

Any bad points entered during the measurement of a track could be erased by the "*" character. When this was detected by REAP, the last pair of coordinates/^{to} have been read from the tape was discarded. By successive use of the "*" character several points along a track could be erased, although in such cases the operator would usually prefer to remeasure the whole track. Incorrect measurements of specific points, for example the apex or a fiducial mark, were not erased using "*", but were remeasured. If REAP encountered more than one measurement of any point or track, then the last rather than the first measurement was accepted. This allowed the operator to repeat any measurements as required.

Although it was quicker and easier to complete the measurement of one view before starting on another, it was possible to move from one view to another if necessary. All that was required was that each change of view was indicated by a view starting character followed by the view number. This facility was found to be useful in cases where incorrect measurements on one view became apparent during the measurement of another view, although this did not occur very often.

2.2.3 The completion of an event

On the completion of an event the ">" character was punched. This was an indication to REAP that there were no further measurements for the event.

Failure to enter the closing bracket would have caused REAP to discard the event. Occasionally it was found that after starting to measure an event it proved to be unmeasurable. By omitting the closing bracket the incomplete event was not passed on to THRESH for further processing.

2.3 REAP

The paper tape output from the measuring machines was read on the IBM 1130 system by the editing and sorting programme REAP. The coordinate measurements and labels were decoded and the data for each event in turn were output onto cards in the format required for input to the geometry programme THRESH.

At the time this was the only way in which data could be input to NUMAC. Since the transmission system from the Durham terminal computer could not handle large amounts of data on account of the slow line rate, the data cards were transported to Newcastle for direct entry to NUMAC. It was not until later that the transmission from the 1130 disk storage became possible.

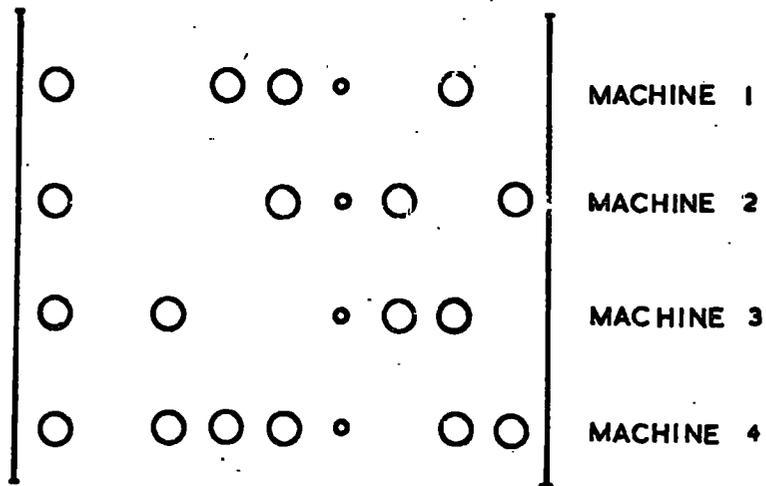
2.3.1 Interpretation of different coordinate codes

At first three different versions of REAP were required to interpret

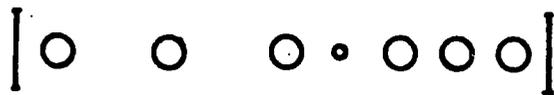
the binary, octal and reflected decimal coordinate codes used by the four machines. As a different colour of paper tape was used on each of the machines, the computer operator could easily determine which version of REAP to use for each tape. At that time, the IBM 1130 computer had only 4K words of core, and it was not until the disk monitor system with 8K words in core was available that it became possible to combine these three versions of REAP into one.

In order that the single version of REAP should be able to interpret the different coordinate codes correctly, it had to be able to determine on which machine each event had been measured. Originally each machine had punched the same event starting character when the operator typed in the opening bracket "<" at the start of each event, but three of the machines were modified so that a different character was punched by each machine. These distinctive characters, which are shown in Figure 2.7, enabled REAP to identify the machine on which each event had been measured, and so to use the correct routine for interpreting the coordinate characters.

The alternative to this was to have used the machine number which appeared in the heading information for each event. This number however, having been entered by the operator, was subject to error and so was not sufficiently reliable to serve as an indicator for REAP. Figure 2.7 also illustrates the coordinate starting character and event terminating character which were used by all four machines.



A. EVENT STARTING CHARACTERS



B. COORDINATE STARTING CHARACTER



C. EVENT TERMINATING CHARACTER

FIGURE 2.7 - SPECIAL CHARACTERS -

2.3.2 Detection of tape punching errors

Besides editing and sorting the data the programme REAP attempted to detect mispunching of the tape whenever possible. Little could be done to correct for mistakes by the operator, such as wrong or omitted labels, but errors of the punching machines could quite often be detected.

Not infrequently two, or even three, coordinate starting characters were punched instead of one. Originally the starting character was not distinguishable from the 6, 10 or 12 coordinate characters of the different machines, and it was not possible to correct for this kind of error. So the machines were modified to give a coordinate starting character which could not be confused with the coordinate characters themselves. It was then possible to detect the double punching of the coordinate starting character, and to eliminate the errors which were caused by this fault.

The distinctive coordinate starting character also made it possible to check that the correct number of characters had been punched for each pair of coordinates. If it was discovered that fewer than the expected 6, 10 or 12 characters had been punched, then that coordinate was discarded. It was found that when too many characters appeared there was usually a trailing blank. In such cases the trailing blank was ignored and the coordinate was accepted.

Sometimes the incorrect punching of the coordinate measurements of a track had the effect of causing REAP to attempt to interpret the last two characters of a pair of coordinates as a point or track label. This could have had serious results as not only would the subsequent coordinates for that track no longer be associated with that track, but they could

be associated with a non-existent track with a nonsensical label. However, by restricting the point and track labels that REAP would accept to those that were actually used, REAP was able to recognise such nonsensical labels and so, instead of associating the remaining coordinates with the nonsensical label, it was able to restore these coordinates to the track to which they belonged. Usually this resulted in the loss of only one pair of coordinates, that pair in which the punching error had occurred.

2.3.3 REAP output

The output from REAP was punched onto cards in the format required for input to the geometry programme THRESH. About 30 cards were required for each event. The last eight columns on each card were not required for data, and so a serial number was punched in the last two columns of each card to identify its position within the event. The cards were then taken to Newcastle where the data were entered directly into NUMAC.

The handling of such large numbers of cards gave rise to certain problems. As with paper tape mispunching could occur. Sometimes this could be detected by the use of checksums. Another problem which always presents itself when cards are handled in large numbers is that of preventing the cards from getting out of order. A programme WEED was written to check the sequence numbers of the cards in each event, and to rearrange the card images if necessary, before the event was passed on to THRESH.

When the NUMAC software was developed to accept programmes and data

not only from cards but from 1130 disk storage, REAP was modified to output the data onto disk instead of onto cards. This reduced by a factor of three the execution time required by REAP and eliminated the problems of mispunching and of cards getting out of order. It was also found to be much more convenient to use disks than cards since one disk can accommodate 200 or more events which would have required about 6000 cards.

2.4 Problems with the off-line system

At first sight the off-line system appeared to provide a satisfactory and flexible method of film measurement. It allowed a wide variety of events to be measured and the measurements of each event could be made in any order. It was also convenient for the operators. Heading information and point and track labels were simply typed in through the input writers. Actual measuring was made easy by allowing the operators to erase measurements, or to remeasure any part of an event. The printed lists on the input writers of all the measurements made enabled the operators to check on the correct typing in of labels, and that all the required measurements were made. The paper tape provided a hard copy of the measurements which could then be reprocessed at a later date if required.

There were only a few restrictions on the system which were imposed for the sake of detecting errors. For example, only a limited set of alphanumeric characters would be accepted by REAP as valid labels. When faults had occurred during the punching of coordinates it was possible for REAP to attempt to interpret certain coordinate characters as point

or track labels. By insisting upon a limited set of permitted labels it was possible to detect such errors, and in some cases only one or two coordinate pairs were lost. This restriction could be lifted if required for any special set of measurements, as it only required a slight modification of the programme REAP.

Although the operators were free to make the measurements for each event in any order, they seldom took advantage of this. Each operator very quickly developed her own order of measuring from which she seldom deviated. It was noticed that in spite of the complete freedom that was allowed, all the operators chose to make their measurements in essentially the same order. This was very satisfactory, as an operator was less likely to omit any of the measurements for an event while following the same sequence of measurements for each event, than by adopting a different sequence each time.

Throughout her shift the operator could also make certain checks on the performance of her machine. By resetting on the apex of the event several times during the measurement of each view she could, by scanning the illuminated coordinate display, check on the correct functioning of the counters. She could also detect the change in rhythm which occurred when the tape punching machine failed to punch the correct number of characters for a coordinate pair. Finally she could check the tape for even parity punching.

From this description it would appear that the paper tape system of measuring was very satisfactory. It was flexible and simple enough so

that the operators could learn to manage the system quite quickly, and it could be used for all kinds of events including special measurements. It was also possible to make various checks to confirm that the data were being transferred to the paper tape correctly.

However, a closer look at the system will reveal that there were some quite serious disadvantages which could not be ignored.

One of the greatest problems was concerned with careless and inaccurate measurements. Usually these were not detected until after the data had been processed by REAP, THRESH and GRIND, which was often some weeks after the measurements had been made. With no immediate evaluation of the accuracy of her work it was all too easy for an operator to attempt to measure as quickly as possible with careless measurements as the result.

By the time the THRESH or GRIND output had been checked, another film was usually on the machine, and it was not until the original film had been put back on the machine that the events which had failed could be remeasured. The whole event was then remeasured in each of the three views, and data were passed on to REAP, THRESH and GRIND in the usual way. This was a lot of work to recover an event which might have failed on only one track.

Although some machine faults could be detected by the operators there were others which were not obvious at the time of measuring, and which were only detected when it was found that THRESH was unable to make a satisfactory reconstruction of the events. As with the events which failed because of careless measurement, so also the events which failed because

of mechanical faults had to be remeasured. In this case it sometimes happened that a large number of events had failed because measuring had continued on a faulty machine before the fault was detected.

The editing and sorting programme REAP checked for, and corrected, as many tape punching errors as possible, but not all of these errors were of such a nature that they could be detected at this stage. So once again there were errors which remained undetected until the output from THRESH and GRIND was available for checking.

There were also minor problems concerned with events which failed because some measurements had been omitted, or because labels had been omitted or entered incorrectly. Errors of this kind should not have occurred very often, as the operators could check each event on the list on the input writer before terminating it with the closing bracket. However, there was no means of ensuring that the operators checked these lists carefully on the completion of each event.

So long as the data were punched onto paper tape and there was no immediate evaluation of the accuracy of the measurements, little could be done about the various problems outlined above. It was with a view to overcoming some of these problems that it was decided to develop an on-line system of measuring.

CHAPTER 3

THE ON-LINE MEASURING SYSTEM

In the on-line system the output from the measuring machines is transferred directly into a computer where the measurements are checked by the programme before the operator is allowed to proceed. Once the machines are connected on-line it is possible to overcome many of the problems that were encountered in the paper tape system. Clearly all errors due to mispunching or misreading of paper tape are eliminated, and the immediate evaluation of the accuracy of the measurements makes it possible for the computer to detect careless measurements and machine faults as they occur.

It also becomes possible for the computer to relieve the operator of some aspects of the routine measuring cycle. In particular, if the computer is first given the topology of the event, it can generate all the necessary track and point labels so that these do not have to be entered by the operator. This requires a certain amount of "conversation" between the computer and the operator, as the operator must know what measurement she is expected to make at any time, and she must be able to signal the computer whenever she is ready to enter a coordinate pair.

The actual form of this "conversation", and the extent to which the data can be checked by the computer, depend very much on the computing system chosen to control the on-line measuring system. The choice of computer for the Durham system was based on several considerations which are discussed in Sections 3.1 and 3.2. The chapter then goes on to

describe the complete on-line system and how it is operated.

3.1 General considerations

Any measuring system, whether off-line or on-line, has to fulfil certain basic requirements. First of all each event that is measured has to be identified uniquely. An eight digit decimal number is found to be convenient for this. Certain other information identifying the type of event, the operator, the date and the machine number is desirable, but not essential.

The measurement of each event requires the recording of 100 or more coordinate pairs. These have to be associated with the points or tracks to which they belong, and with the appropriate views. The labelling of the points and tracks should not only identify them, but should also indicate any special properties. Stopping tracks, for instance should be distinguished from those that leave the chamber.

It is also important that the system should be able to handle a wide variety of different events, and that some facility for erasing and remeasuring points should be included.

Complete versatility would accommodate an infinite variety of events, making no restriction on the number of points or tracks that could be measured for any one event. Neither would there be any restriction on the total number of coordinate pairs that could be entered, or on the choice of alphanumeric characters which could be used for labels. In practice, however, only a limited number of points, tracks and identifying labels

are required for each event, and nearly all events fall within a certain basic pattern. With few exceptions each event is measured in three views. Measurements in each view are made of the four fiducials, the apex, other points such as stopping tracks or "kinks", the beam track, and a variable number of secondary tracks some of which may have special properties which should be indicated by special labels.

These requirements are fundamental to any successful measuring system and were all fulfilled by the paper tape system which has been described in Chapter 2. An on-line system, however, can provide a better method of measuring simply because the measuring machines are in direct contact with a computer throughout the measuring process.

By coding each event according to its topology, and entering this coded topology into the computer, it is possible for the computer to generate a list of the required measurements and the corresponding labels. The operator can then proceed to measure the event according to this list without having to type in labels to identify each measurement. This eliminates failures of events due to measurements being omitted, or to incorrect labels being entered by the operator.

The measurement of each event then comes under the control of the computer, rather than the operator, which requires that the measurements for each event are made in a fixed order. This is not such a restriction as might at first be supposed, as even under the paper tape system which, as already mentioned, allowed complete versatility in this respect, each operator developed her own order of measuring and seldom deviated from it.

The computer then has to be able to give certain commands to the operator to indicate what measurements are required, and, after receiving these measurements, to indicate whether or not they were acceptable. Similarly the operator must be able to give instructions to the computer if she is to be free to make erasures or to remeasure any part of an event. Some facility must also be provided to enable her to enter the necessary heading information which identifies each event.

All these general requirements had to be considered in making the choice of the computer for the on-line system. The other considerations that influenced this choice are discussed in the next section.

3.2 Choice of computer for the on-line system

Some of the general considerations for ^{an} on-line system have already been discussed in the previous section. There the emphasis was on the control of the measuring machines and the input of data to the computer. Another requirement is that the data should be output in a form suitable for further processing.

Most of this further processing is carried out in NUMAC which, it is relevant to recall, comprises an IBM 360/67 computer situated in Newcastle, with an IBM 1130 terminal computer in Durham. For the on-line system to be compatible with NUMAC the output must be in one of two forms of input to NUMAC. The first of these is by direct entry to NUMAC from 9-track magnetic tape. The second is by remote entry to NUMAC from magnetic disk on the terminal computer in Durham.

The limited financial backing restricted the choice of computer to those having a core storage of about 8000 words and a minimum of input and output devices. Two systems were considered which would have fulfilled the necessary conditions. These were the DEC PDP8 system, and the IBM 1130 system.

The PDP8 system had some advantages in that it was slightly the cheaper of the two for a similar configuration. It also had a faster core cycle time of 1.5 μ sec compared with 3.6 μ sec of the 1130 system. However, the IBM 1130 system was chosen for the following reasons.

Output from the PDP8 would have been onto a small temporary disk, and from there it would have been transferred to magnetic tape at the completion of each event. About once a week these tapes would have had to be taken to Newcastle for direct input to NUMAC. With the 1130 system the data are output directly onto magnetic disk. Each day the disk is taken to the Durham terminal computer, some 200 yards away, for remote entry of the data to NUMAC. This latter arrangement is clearly more convenient than the former.

As it was planned to have up to six measuring machines connected on-line to the computer it was important, with only 8000 words of core, to consider how efficiently the core space could be used. In the PDP8 system the basic word length is only 12 bits so that double precision would have been needed to handle the 16-bit coordinate measurements from the measuring machines, and three or four words would have been required to store each coordinate pair. The 1130 computer on the other hand, has a

basic word length of 16 bits, and so requires only two words to store each coordinate pair. Furthermore, most 1130 programmes require fewer words of core than the corresponding PDP8 programmes because the 1130 computer has 29 basic instructions compared with only 12 of the PDP8 system.

In either case most of the programmes would have to be written in Assembler language or its equivalent, as only in this way would it be possible to keep the core requirements within the 8000 or so words that would be available. Now there was already available in the University a considerable amount of 1130 programming experience. This experience had been gained on the Durham terminal computer for NUMAC, and the author herself had gained some valuable experience of 1130 Assembler language programming when she wrote the new version of REAP which was required for the paper tape system. No such experience of PDP8 programming was available and so the writing of Assembler language programmes for the PDP8 computer would have been much more difficult than it was for the 1130 computer.

Another advantage of the 1130 was that it would be possible to do most of the programme development on the Durham terminal computer for NUMAC before the on-line system was installed. No PDP8 system was available, so that if that had been chosen, the programme development would have had to have waited until the on-line system had been installed. As already mentioned, this system would have a minimum of input and output devices and so programme development would have been slow using paper tape, with no line printer, and with an unfamiliar system. It would probably not have been until about a year after its installation that an on-line system with a

PDP8 computer would have been operational.

With the 1130 system however the programme development could all be done beforehand using the Durham terminal computer for NUMAC. The importance of this computer was not confined to its being available before the installation of the on-line system, but was due also to the fact that it has both a card reader and a line printer. This made programme development much easier than it would have been using the on-line computer which has neither of these. Consequently it was possible to compile and test nearly all the programmes for the on-line system before that system was installed, with the result that the on-line system was operational within a few weeks of its installation.

3.3 The on-line measuring system

The on-line measuring system is based upon an IBM 1130 computer. The measuring machines are connected to the computer through a Westinghouse interface built specially for the purpose, and an OEM channel. Each machine has a control panel through which the operator can communicate with the computer, and vice versa.

Further communication between the computer and the operator is achieved by relaying printed messages by closed circuit television to individual monitors placed beside each measuring machine.

Measurements from each machine are checked by the computer before the operator is allowed to proceed. Good measurements are output onto magnetic

disk. The data are later sorted and edited before being submitted to NUMAC for further processing.

Provision is also made for another experiment to run on-line to the computer. This experiment, on cosmic ray research, involves the reading in of data from a Mullard core store. No on-line data processing is required, only the transfer of the data onto a magnetic disk. A second disk drive was installed so that the data from the core store of the cosmic ray experiment could be output onto^a separate disk from that used for the data from the measuring machines. The core store is connected to the computer through the Westinghouse interface and the OEM channel, and may be run at the same time as the measuring machines. Some extra circuitry in the interface was required to transfer the data from the core store to the computer.

The complete system can accommodate six measuring machines and the core store all operating simultaneously.

3.3.1 The computer

The Central Processing Unit (1131 CPU) of the IBM 1130 computer has a basic word length of 16 bits and has 29 basic instructions. The model used for the on-line control of the measuring machines has 8192 (8K) words of core, and has a core cycle time of 3.6 μ sec. A faster version is available with a core cycle time of 2.2 μ sec, but this is considerably more expensive, and the slower version is sufficient for this application. It would also have been possible to obtain the 1130 computer with 4K, 16K or 32K words of core, but a core of 4K words would have been inadequate

for this application, and financial support was only available for an 8K system.

Built into the basic unit of the computer is the console printer which can be used for any printed output required, the console keyboard which can be used to enter instructions or data into the computer, and the main disk drive which takes the disk cartridges. Each disk has space on it for 500K words, of which about one third are usually taken up by the Disk Monitor System, the Fortran and Assembler Compilers and various programmes. A second disk can be mounted on the additional disk drive which requires a multiplexer for connection to the computer. The system also includes a paper tape reader, a paper tape punch, and an OEM channel through which the interface is connected to the computer. The complete system is illustrated in Table 3.1.

Table 3.1

Durham On-Line Computing System

UNIT NUMBER	DEVICE
1131	Central Processing Unit
1134	Paper Tape Reader
1055	Paper Tape Punch
2310	Additional Disk Drive
1133	Multiplexer for 2310
2315	Disk Cartridge

Except for the second disk drive this is little more than the minimum configuration for an 1130 system. It is adequate for the on-line control of the measuring machines as it provides for the collection of data directly from the machines and the output of data onto disk. The console printer is used for printing coded error messages and for providing a brief log of the measurements made.

3.3.2 The measuring machines

Only minor changes had to be made to the measuring machines which had been used for off-line work to incorporate them into the on-line measuring system.

(a) The coordinate measurements

One of the main differences between the on-line system and the paper tape system is that the coordinate measurements are no longer transferred to paper tape punching machines, but are connected to the interface for direct transfer to the computer. The basic word length of the computer is 16 bits and transfer of data to and from the computer is also in units of 16-bit words. All information has to be presented in this way before it can be transferred through the interface. In particular the coordinate measurements have to be made available in binary patterns of 16 bits.

For machines 1 and 2 this was not very difficult as the coordinate measurements were already in binary form having been formed in 17-bit binary counters (see Section 2.1.1). Now the X coordinate represents a measurement along the length of the film and all 17 bits of the counters

are required to accommodate the corresponding range of values. However, the accuracy with which an operator can set on any point is of the order of 4 fringes. So there will be no real loss of precision if the least significant bit in the counters is dropped, while the other 16 bits are transferred to the computer. Since all the beam tracks and the majority of secondary tracks are roughly parallel to the X axis, any loss of precision in the X direction will have little effect on the calculated value of the sagitta.

On the other hand, the accuracy with which the Y coordinates can be measured is important as this determines the accuracy of the calculated sagitta for most tracks. Since the Y coordinate represents a measurement across the width of the film the full precision can be preserved as only 16 bits are required to accommodate the smaller range of values. In this case there is no loss of precision at all since the most significant bit of the 17 in the counter is neither used nor transmitted.

So the coordinate measurements prepared by machines 1 and 2 can easily be made available in the required form for input to the computer. The least significant of the 17 bits in the X counters is not transmitted, and the most significant of the 17 bits in the Y counters is neither used nor transmitted. This effectively doubles the fringe size in the X direction, but the full precision is preserved in the Y direction where it is of greater importance.

Machine 3 however, required a more fundamental modification to be

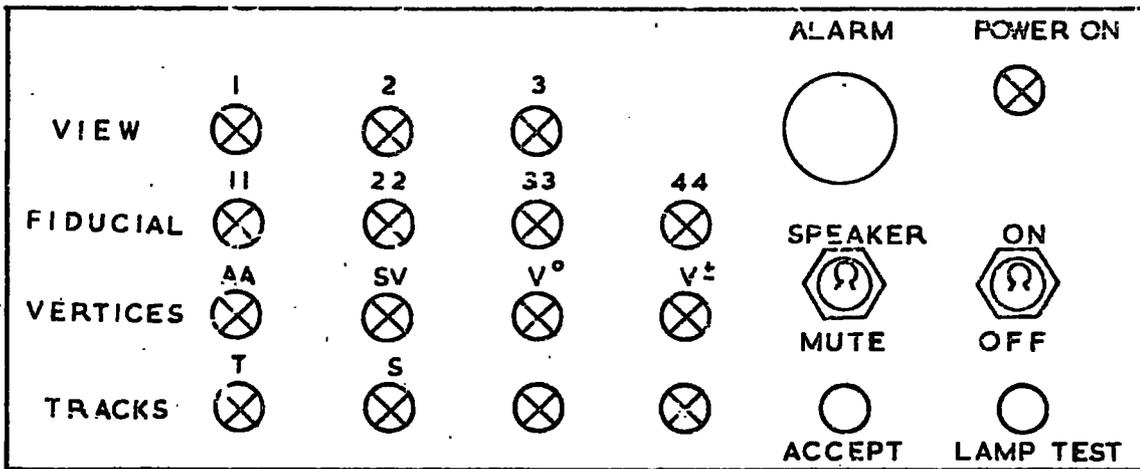
made. It has already been mentioned (see Section 2.1.2) that the digitisers on this machine produced coordinate measurements in a reflected decimal code. To continue to use these digitisers would have meant building circuitry to convert the reflected decimal coordinates into the required binary form. At the time it was found to be cheaper in both time and money to replace the mechanical digiters with rotary Moiré fringe digitisers which produce the coordinate measurements in binary form as required. The same considerations of precision and of the required range of values for X and Y coordinates apply to this machine as they did to machines 1 and 2, so that the effective fringe size in the X direction is again twice that in the Y direction.

The octal coordinates of machine 4 were easily converted to binary as each octal digit translates directly into three binary digits. Once again the ranges of values required for the X and Y measurements are such that the effective fringe size in the X direction is made twice that in the Y direction.

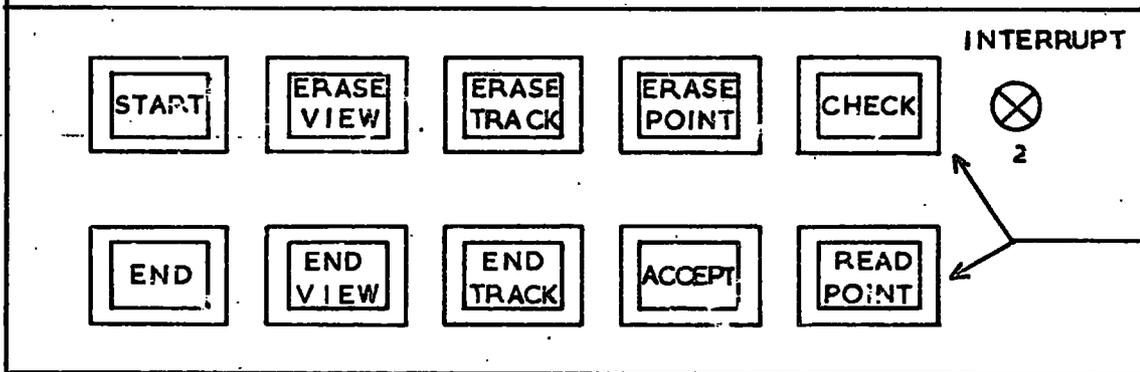
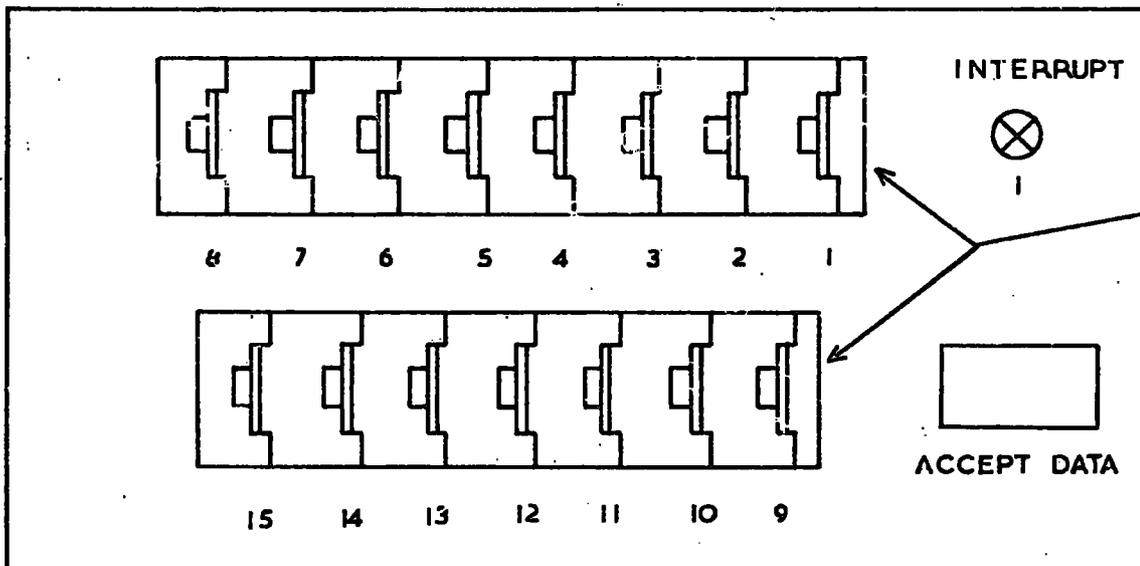
(b) The control panels

Another significant alteration to the measuring machines is that the input writers for entering heading information and labels are no longer required. Instead, each machine has a control panel which enables both the operator to communicate with the computer, and the computer to communicate with the operator (see Frontispiece).

Each panel consists essentially of two sections which are illustrated in Figure 3.1. The first section consists of a lamp display which is an array of 15 lamps which may be illuminated or extinguished by signals



SECTION 1 LAMP DISPLAY



SECTION 2 SWITCH PANEL

FIGURE 3.1 — OPERATOR'S CONTROL PANEL —

from the computer. Alongside the lamps is an alarm which can be turned on by the computer but which can only be turned off by the operator. Together the 15 lamps and the alarm correspond to a signal 16-bit binary word, and they are controlled from the computer by the output of a 16-bit digital output word. The pattern of illuminated lamps indicates to the operator what measurement she is required to make at any particular time. The alarm is used as a signal to the operator that her latest measurements are not acceptable, or that she has made some kind of error.

The second section of the control panel contains 15 decade edge switches, and a set of 10 latching illuminated switches (see Figure 3.1). These are set by the operator and are read by the programme whenever the operator requests attention by pressing one of her interrupt buttons.

In this way the edge switches are used to enter the heading information for each event which, as in the off-line system, consists of the event number, the experiment number, the operator number, the date and the topology. In this case the machine number does not need to be entered explicitly.

The latching switches, which are illuminated when set in the "on" position, enable the operator to indicate to the computer what she wants to do. Most often this is to enter a pair of coordinate measurements. Other requests, such as to erase the last point or track, or to indicate that the measurement of a track has been completed, can all be made through the selection of different switches.

The operator requests the attention of the computer by means of two interrupt buttons. The first of these, called ACCEPT DATA, is mounted on the control panel and has the special function of indicating the start of an event. The other, called EXECUTE, is mounted close to the controls of the measuring machine and is used in all cases other than for the start of a new event. The actual function required depends upon which of the illuminated switches on the control panel are on. Unless at least one of these switches is on, the EXECUTE interrupt button is inhibited.

Two interrupt indicator lamps are mounted on each control panel. Each time an interrupt is raised, the corresponding lamp on the control panel of the machine concerned goes out. It comes on again when the computer responds to the interrupt. The response is usually so rapid that only a slight flicker of the lamp can be detected.

3.3.3 The OEM Channel

The measuring machines are linked to the computer through the Westinghouse interface and the OEM channel.

The OEM channel, which is a standard IBM fitting, connects directly with the CPU and is designed to provide for the attachment of a variety of devices. It consists essentially of:

- (i) two 16-bit registers, one each for input and output,
- (ii) a single 4-bit register for control,
- (iii) two lines of data register status for synchronising,
- (iv) four lines of device status sense, and
- (v) a 6-bit register for interrupt requests.

All movement of data through the OEM channel is under direct programme control from the CPU. The input register can be read into core. Data from any word in core can be transferred into the output register. The four bits in the control register can be set or reset by the programme. The two data register status lines control the movement of data to and from the input and output registers. The device status lines and interrupt request register may be sensed by the programme.

Devices connected to the OEM channel can request the attention of the computer through the interrupt request register. The setting of any of the six bits in this register causes a programme interrupt to be raised in the CPU. The programme then determines the cause of the interrupt by examining the interrupt request register to see which bits have been set.

3.3.4 The Westinghouse Interface

An interface was required to connect the measuring machines to the OEM channel. The main functions of the interface are:

- (1) to set the appropriate bit in the interrupt request register of the OEM channel whenever an interrupt button on one of the measuring machines is pressed,
- (2) to fetch data words from the measuring machines and transfer them to the input register of the OEM channel from where they may be read into the computer,
- (3) to route digital output words from the output register of the OEM channel to the lamp display of a specified machine.

Besides these basic functions the interface must be able to inform the computer which machine raised an interrupt and which of its two interrupt buttons had been used. This requires a register within the interface which has at least 12 bits, one corresponding to each of the 12 interrupt buttons. The interface has also to inhibit each interrupt button when it has been used to raise an interrupt. The interrupt button is only released when the computer has responded to the interrupt in a prescribed way.

In addition the interface is required to handle the transfer of data from the core store of the cosmic ray experiment to the input register of the OEM channel.

Clearly such an interface is not a standard piece of equipment, but had to be designed and built specially for the purpose. This task was undertaken by Westinghouse Brake and Signal Company, a firm which already had some experience of building interfaces for the IBM 1130 computer. This firm also built the control panels for the measuring machines which have been described in Section 3.3.2.

The author was closely involved in the discussions with Westinghouse over the design of the interface, and in particular over the way in which it would be handled by the programme in the CPU.

Only a brief description of the interface is given here as full details of its operation will be given in the next chapter.

(a) Interrupts

The raising of an interrupt requires a bit to be set in the interrupt

register of the OEM channel. In order to distinguish between interrupts from the measuring machines and interrupts from the core store of the cosmic ray experiment, two different bits in the interrupt register are used, depending on the cause of the interrupt. This makes it possible for the computer to determine the cause of each interrupt since it can read and examine the interrupt register.

(b) Input

Since the interface is connected to the computer through the OEM channel all movement of data is controlled by the programme in the CPU. The mode of operation, whether input or output, is governed by the different settings of the four control lines.

There are seven input data words associated with each of the six measuring machines. These are identified within the interface by six sets of seven consecutive addresses. For each machine the first four addresses refer to the 15 decade edge switches, the fifth refers to the settings of the 10 latching switches, and the last two refer to the X and Y coordinate registers respectively. A total of 42 input addresses are required for the six measuring machines. There are also two other input addresses within the interface which refer to a device status word, and to the core store of the cosmic ray experiment respectively.

The programme specifies which data word is required at any time by first sending the address of that word to the interface, which will then present the contents of that address on the input register of the OEM channel from where it may be read into the CPU by the programme. About

20 μ sec elapses after the address has been given before an interrupt is raised to inform the programme that the data word is available.

(c) Output

A similar procedure takes place in the output mode. A unique address is assigned to each of the six control panels. The programme first gives the address of the control panel required, so that when the digital output word is presented by the OEM channel it is routed by the interface to the required lamp display. Of the 16 bits of the output word, 15 are used to control the lamps, and one may be used to sound the alarm.

(d) Westinghouse device status word

When an operator requests the attention of the computer by pressing one of the interrupt buttons on her machine, the computer must be able to determine which machine raised the interrupt. A register within the interface, called the Westinghouse device status word, holds this information. This word can be read by the programme in the same way as data words. Two bits are associated with each machine, one for the ACCEPT DATA interrupt, and one for the EXECUTE interrupt, which are set when the corresponding buttons are pressed. By examining this word the programme can determine which machine raised the interrupt, and whether it was ACCEPT DATA or EXECUTE.

Having determined which machine requires attention, the programme can then read in the seven data words from that machine. As the fourth data word is read, the ACCEPT DATA bit within the device status word is reset. The EXECUTE bit is reset as the Y coordinate is read in. Usually

only one of these two bits will have been set at any time.

(e) Monitoring

In order to facilitate the testing of the transfer of data through the interface a monitoring panel has been built into its main control enclosure. Incoming data can be displayed on an array of 16 lamps. Outgoing data can be simulated on an array of 16 switches. A further three switches are provided which can be used to set or reset control lines 1, 2 and 3.

By means of these switches, and an initiate button which is also on the monitoring panel, the programme instructions which are required to control the movement of data through the interface, can be simulated. The array of lamps provides a quick and easy check on the correct input of data. Output of data can be checked by glancing at the lamps on one of the control panels.

The panel also includes several monitoring points which may be used for dynamic testing of the interface.

During the initial testing of the on-line system the monitoring panel was particularly useful as it provided a simple way of checking the correct transfer of data through the interface. It is also useful in that, should the system break down at any time it may be used to determine whether or not the interface is operating correctly. Since two manufacturers are involved in the design and construction of the on-line system it is desirable when a breakdown occurs, to be able to determine whether the fault lies in the IBM equipment, or in the Westing-

house part of the system. The purpose of the monitoring panel is to make it relatively easy to determine the location of the fault. However, the whole system has proved to be so reliable that it has never been necessary to make such a diagnosis.

The monitoring panel is seldom used except for the AUTO/MANUAL switch, which it is found convenient to set to MANUAL whenever the computer is being operated on its own without the measuring machines. Accidental interrupts from the measuring machines are then inhibited and so cannot interfere with the running of the computer.

3.3.5 Closed circuit television

The system also includes closed circuit television. The camera is positioned over the console printer, and a monitor is placed close to each measuring machine. This enables each operator to read any messages that may appear on the console printer without getting up from her machine.

In this way the one console printer may be used for giving supplementary information to all the operators. Provided no errors are made, only the event numbers are printed, but coded error messages tagged with the relevant machine numbers are printed whenever failures or illegal entries are made. The computer draws the operator's attention to an error message by sounding the alarm on her control panel. In the case of failures, the printed message gives some indication of the magnitude of the failure. The printer can also be used to give in more detail the request for the next measurement which appears on the lamp display of the control panel.

3.4 The operation of the measuring machines

Before starting to make any measurements an operator must first sign on. The experiment number, operator number and date are read into the computer from the edge switches on the control panel, and are used to form the heading information for all the events measured on that machine until a new operator signs on.

The start of an event is indicated by the ACCEPT DATA interrupt button. The event number and topology are read from the edge switches. The computer decodes the topology, generates a list of the required measurements and the corresponding labels, and then calls for the first measurement by displaying the appropriate pattern of lights on the lamp display.

The operator proceeds to measure the points and tracks of the event in the order requested by the computer as indicated by the lamp display. Most of the measurements are checked before the operator is allowed to proceed. When measurements are not acceptable the alarm is used to draw the operator's attention to the failure, and a remeasurement is requested.

At the completion of each view the operator has to reset on the apex of the event to check on the correct functioning of the counters on the machine. The event is completed at the end of the third view. All the lights on the lamp display come on to indicate that the computer is ready to receive the start of the next event.

Before leaving the machine at the end of her shift the operator should sign off, so that the next operator to use the machine may sign on.

3.4.1 Signing on

The computer will not accept the start of an event unless the operator has signed on. At the beginning of her shift she must therefore carry out the following procedure. Firstly she should set the START switch on the control panel. She should then set the experiment number, her the operator number and /date on the decade edge switches as shown in Figure 3.2A. Finally she should press the EXECUTE interrupt button.

This interrupts the programme in the CPU. In response to the operator's request for attention, the computer reads in all seven data words from her machine. The setting of the START switch tells the computer to interpret the edge switch settings as experiment number, operator number and date, and to store these numbers in core so that they may be used for the heading information for the subsequent events that will be measured on that machine.

The computer then causes all the lamps on the display panel to be illuminated so that the operator knows that she may proceed with the start of her first event.

3.4.2 The start of an event

Some of the heading information required for each event is read into the computer when the operator signs on. The remaining items, the event number and topology coding, are entered when the measurement of the event is started.

EXPERIMENT NUMBER	8 0 0 0 0 0 3 1								⊗	OPERATOR NUMBER
	0 1 8 0 3 7 1									
										DATE
	START								⊗	

A - SIGNING ON -

EVENT NUMBER	3 6 5 0 7 2 8 1								⊗	
	4 0 4 0 0 0 0									
TOPOLOGY CODING										ACCEPT DATA
									⊗	

B - THE START OF AN EVENT -

NOTE - THE MEASUREMENTS FOR THIS EVENT ARE MARKED
* IN TABLE 3.2

FIGURE 3.2 - EXAMPLES OF CONTROL PANEL SETTINGS -

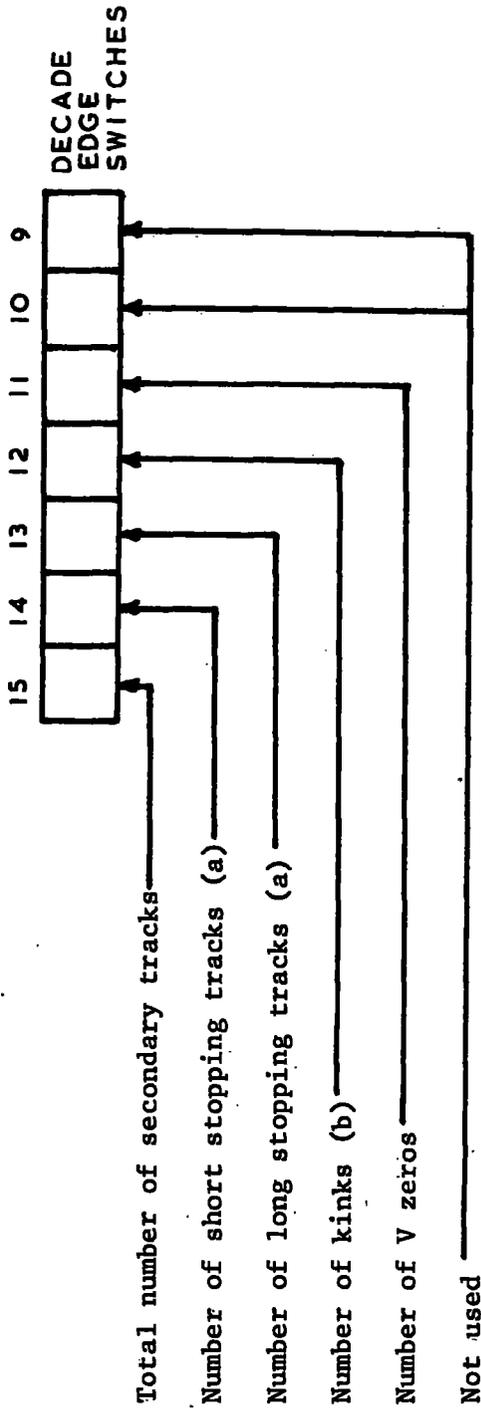
The event number is the same as in the paper tape system, consisting of an eight digit number made up of the film number, the frame number and the event number. The topology however is different as it has to give much more information than it did in the paper tape system. Seven digits are available which can be used to give the computer a full description of the event, indicating the total number of secondary tracks, the number of pion and proton stopping tracks and whether they are "short" or "long". In this respect a "short" track is one whose end point only is to be measured. For a "long" track both the end point and at least five points along the track will be measured. The topology also indicates any kinks or V zero tracks that are to be measured. The coding of the topology is illustrated in Figure 3.3.

To start the event the operator should set the event number and topology coding on the edge switches as shown in Figure 3.2B, and then press the ACCEPT DATA interrupt button. This causes a programme interrupt in the same way as the EXECUTE button. The computer responds as it does to any interrupt by reading the seven data words from the machine.

The topology is checked against certain limitations, some of which depend upon the space reserved in core for the generated labels, and some of which check the legality of the topology. Provided the topology is acceptable the heading information for the event is moved to the disk output area from where it will be transferred to the disk, and a list of all the measurements required for the event is prepared. These measurements are selected from the list shown in Table 3.2. The computer then calls

Figure 3.3

Topology Coding



(a) Count 4 for a proton, 1 for a pion

(b) Count 4 for a long kink, 1 for a short kink

Table 3.2

List of measurements and corresponding lamp displays

Note: The appropriate view light is also illuminated

NUMBER	LABEL	LAMP DISPLAY						DESCRIPTION						
		TRACKS		VERTICES		FIDUCIAL			VIEW	A				
		S	T	AA	SV	V ^o	V ⁻		44	33	22	11	1	2
1	*AP			1										Apex
2	*A9	1		1										Beam track
3	*11											1		Entry fiducials
4	*44							1						
5	XX			1										
6	WW			1										Short stopping protons
7	YY			1										
8	ZZ			1										Short stopping pions
9	*PP			1	1									
10	*AP	1		1	1									Long stopping protons
11	OO			1	1									
12	AO	1		1	1									
13	QQ			1	1									
14	AQ	1		1	1									Long stopping pions
15	RR			1	1									
16	AR	1		1	1									
17	UU					1								
18	U1	1				1								
19	VV					1								
20	V1	1				1								Short kinks
21	AT	1		1	1									
22	TT			1	1									
23	T1	1				1								
24	AS	1		1	1									Long kinks
25	SS			1	1									
26	S1	1				1								
27	*A1	1		1										
28	*A2	1		1										
29	*A3	1		1										
30	A4	1		1										
31	A5	1		1										
32	A6	1		1										
33	A7	1		1										
34	A8	1		1										
35	MM			1										
36	M1	1		1										
37	M2	1		1										
38	NN			1										
39	N1	1		1										V zeros
40	N2	1		1										
41	*22								1					
42	*33							1						Exit fiducials
43	*AA			1										Apex check

*These measurements are required for the event whose topology coding is shown in Figure 3.2B

for the first measurement of the event by causing the appropriate lights on the lamp display to be illuminated.

Provided the operator has signed on, the ACCEPT DATA interrupt will always be interpreted as the start of a new event. Usually this is given only immediately after signing on, or after the completion of the previous event. If however the ACCEPT DATA button is pressed during the measurement of an event the unfinished event will be abandoned, and the measurement of the new event initiated. When the data are sorted and edited before being input to NUMAC, such an incomplete event will be discarded.

3.4.3 The measurement of a view

Once the measurement of an event has been initiated by the ACCEPT DATA interrupt, and the topology has been accepted by the computer, the operator may proceed to measure the event.

The set of measurements that will be required for the event are selected from the list in Table 3.2 according to the topology that was specified. Whatever topology was given, the selected measurements will always include the apex, the beam track and four fiducial marks. All the measurements that are required for the event will be called for by the programme in the order in which they appear in Table 3.2. This table also illustrates the patterns of lights which are used to request each of the 43 possible measurements.

This measuring sequence is similar to that which was adopted by

most of the operators when they worked with the paper tape system. Each view is started with the measurement of the apex and the beam track, followed by the two entry fiducial marks, and is concluded with the measurement of the two exit fiducial marks and a check on the apex. The coordinates of the entry fiducials are held in core until the exit fiducials are measured, thus making it possible to provide an overall check on the accuracy of the view by examining the relative positions of the fiducials.

The secondary tracks which constitute the rest of the event are also required to be measured in a particular order. Essentially this involves starting close to the apex of the event, and then moving outwards. The special tracks are measured first, starting with short stopping tracks, then the long stopping tracks and kinks. These are followed by ordinary secondary tracks leaving only the V zeros to be measured before the exit fiducials. The same sequence of measurements is repeated in each of the three views.

(a) Apex

At the start of each view the lamp display indicates to the operator that she is required to measure the apex of the event. In the paper tape system the coordinate counters were set at the start of each view with the reference mark of the measuring machine in coincidence with the apex of the event (see Section 2.2.2). A similar procedure is required in the on-line system.

First of all the operator should set the READ POINT switch on her

control panel, and make sure that all the other latching switches are off. She should then bring the apex of the event into coincidence with the reference mark, set the counters and raise an interrupt by pressing the EXECUTE button.

The READ POINT setting on the panel indicates to the programme that a pair of coordinates is available for the last measurement requested, which is shown on the lamp display. In this case this is the apex of the event.

For this particular measurement only one pair of values is acceptable, namely the setting values. With only 16 bits for each coordinate measurement, the setting value for both X and Y is 2^{15} . If this is not the value which is read in for both X and Y, then the computer calls for a remeasurement of the apex. This is done by sounding the alarm, printing a coded message and repeating the same request by illuminating the same pattern of lights on the lamp display.

(b) Beam track

Provided the measurement of the apex is acceptable the operator is requested to proceed with the measurement of the beam track. This is indicated by a new pattern of lights appearing on the lamp display. With the READ POINT switch still on, the operator should measure about nine points along the beam track. Each point is read into the computer in response to an EXECUTE interrupt. After the last point has been read, the operator should indicate the completion of the track by setting the END TRACK switch on the control panel, and then pressing the EXECUTE

button.

The programme then makes a fit to a circle of the measured points of the beam track. Provided this is a good fit (see Section 4.6), the coordinates are moved to the disk output buffer, and the programme calls for the next measurement. If the track measurements are not acceptable, then the operator is requested to remeasure the track. This is done by sounding the alarm, printing a coded error message and repeating the same request on the lamp display.

(c) Entry fiducials

After an acceptable set of coordinates has been received for the beam track the computer requests the measurement of fiducial 11, followed by fiducial 44, the entry fiducial marks. These are entered in the same way as all other points, by pressing the EXECUTE button with only the READ POINT switch set on the control panel.

(d) Secondary tracks

The secondary tracks come next in the sequence of measurements. Any short stopping tracks whose end points only are required, will be called for first. Proton stopping tracks will be requested before pion stopping tracks. If there are two or more of any kind of track to be measured in any event, these should be measured in anticlockwise rotation from the beam track. This is to ensure that the track and point labels are correctly assigned in each of the three views.

In the case of two similar tracks lying close together it is not always possible to determine which corresponds to which in each view.

However, there is a facility in the geometry programme THRESH, called track matching, which can detect when two close secondaries have been confused, and can recombine the measurements from the different views correctly.

After the short stopping tracks, if any, have been measured, the computer will call for any long stopping tracks that were indicated by the topology coding of the event. For each of these the end point is measured first, followed by about nine points along the track. The end of each track is indicated in the same way as for the beam track, and the computer checks the measurements by making the same circle fit calculation. In the case of very curved tracks the test for an acceptable fit is less stringent (see Section 4.6).

Any other special measurements such as kinks or V zeros, and any ordinary secondary tracks, are each called for in turn by the computer. The track test for ordinary secondary tracks includes a check on the sagitta which is required to be at least 2 mm in the bubble chamber.

(e) Exit fiducials

When all the tracks have been measured, the exit fiducials 22 and 33 are called for.

(f) Apex check

Before leaving one view to go on to the next view the operator is requested to return to the apex of the event. The coordinates are read into the computer and are checked against the setting value which was entered at the start of the view. This is to provide a check on the

counters of the machine so that failures of the counters will be detected by the computer. If the check is not acceptable the operator should reset the counters, and will be requested to remeasure the last part of the view during which the failure must have occurred. If the check is acceptable the operator proceeds with the next view.

The end of view 3 marks the completion of the event. Throughout the measurement of the event the coordinate measurements are transferred to the disk as they become available, and the programme keeps a record of which coordinates belong to each track. At the close of the event this referencing information is also transferred to the disk, and all the lamps on the control panel come on to inform the operator that she may start the next event, or sign off if she has come to the end of her shift.

3.4.4 Signing off

Before an operator leaves her machine at the end of her shift she should sign off. If she does not do this the next operator will not be able to sign on.

In order to sign off the operator sets the END switch on her control panel and then presses the EXECUTE interrupt button. Provided she was not in the process of measuring an event, this will be accepted and all the lamps will be turned off.

Signing off during the measurement of an event requires two interrupts from the EXECUTIVE button. The first of these causes the alarm to be sounded as a warning, and the second forces acceptance of the request to

sign off. The unfinished event will be discarded and must be remeasured from the start by the next operator.

3.4.5 Additional Features

Besides being able to control the measurement of an event in the manner which has just been described, the on-line system also provides certain additional features. These enable the operator to request various functions by setting different latching switches on her control panel before pressing the EXECUTE interrupt button.

The special functions of some of the latching switches have already been mentioned.

(a) START

The START switch is used when the operator wishes to sign on at the beginning of her shift (see Section 3.4.1).

(b) END

The Operator uses the END switch when she signs off at the end of her shift (see Section 3.4.4).

(c) READ POINT

During the measurement of an event the READ POINT switch is used to indicate to the computer that a pair of coordinates is available and should be read in (see Section 3.4.3).

(d) END TRACK

When the operator has completed the measurement of a track she sets the END TRACK switch which causes the circle fitting routine to check the coordinates (see Section 3.4.3).

The other switches provide additional facilities which may or may not be required during the measurement of any event. These are described below.

(e) CHECK

The CHECK switch is used when the operator wants more information than is given on the lamp display about the next measurement required of her. Although each point and track in any event has a unique label within that event, there are several instances where the same pattern of lights ^{for} is used to call/more than one particular measurement. For instance, if an event has two ordinary secondary tracks, Table 3.2 shows that their generated labels will be A1 and A2, but they are both called for with the same display of lights. Under normal circumstances this causes no confusion as the operators are familiar with the system and know the order in which the measurements are called for. If however, an operator is interrupted at any time, it may be that, when she returns to measurement, she cannot tell from the lamp display which measurement is required of her. By setting the CHECK switch on the control panel, with all the others off, and pressing the EXECUTE interrupt button, she can determine exactly how much she has done as the computer will then print a message which gives the generated label of the measurement required, and if this is a track label, the number of points already measured on that track.

(f) CHECK/READ POINT

The CHECK switch has another, completely different function when it is used together with the READ POINT switch. It is used to check on

the counters of the machine. The operator should first bring the reference mark of her machine into coincidence with the apex of the event being measured. She should then set both the CHECK and the READ POINT switches, and press the EXECUTE button.

This causes the computer to read the coordinates and compare them with the setting value which is required for the apex at the start of each view. Provided these two values are not more than four fringes apart this is considered acceptable, and the operator is allowed to proceed. If the resetting is not acceptable, then the operator has to reset the counters and remeasure from where they were last considered to be satisfactory. This would be where the last CHECK/READ POINT test was made, or where the last track measurement was accepted, whichever is more recent. An acceptable track measurement can be considered to be a satisfactory apex check because one of the checks made by the curve fitting procedure is that the apex lies close to the curve which is fitted to the measured points.

If the apex check fails, the computer calls for the required measurement by illuminating the appropriate pattern of lights on the lamp display. It also causes the alarm to be sounded, and prints two messages. The first of these indicates the amount by which the coordinates failed to correspond with the correct setting value, and the second gives the label of the measurement from where the operator should start to remeasure.

(g) ERASE POINT

When an operator wishes to delete the last pair of coordinates that were read from her machine, she uses the switch marked ERASE POINT.

If this is followed by the READ POINT setting of the switches then the new pair of coordinates will replace the old pair. If it is followed by the END TRACK setting, then the last point entered will be ignored. The ERASE POINT switch may be used several times in succession to delete several points, but will not be accepted immediately following the completion of a track, or the completion of a view.

(h) ERASE TRACK

If it is required to erase all the points on a track, as would be the case if the operator had been measuring the wrong track, then this can be achieved by setting the ERASE TRACK switch. The subsequent points to be entered will then replace those previously entered, so that the whole track must be remeasured. This setting is allowed during the measurement of a track, or immediately following its completion. It may be used repeatedly to delete more than one track. It is not accepted immediately following the completion of a view or just after an individually labelled point such as the apex, a fiducial or a stopping point, has been measured.

(i) ERASE POINT/ERASE TRACK

The ERASE POINT and ERASE TRACK switches may be used together. In this case the programme determines whether the last measurement was of a point or a track and deletes accordingly. This setting may also be used several times over to step backwards through the view one point or track at a time. Each of the measurements deleted must be remeasured and will be called for in turn in the usual way. This setting will only delete measurements within the current view.

(j) ERASE VIEW

If it is required to remeasure any part of a view which has been completed, then the ERASE VIEW switch must be used. When this is set during the measurement of a view, all the measurements on that view will have to be repeated, but if it is used immediately following the completion of view 1 or view 2 then that completed view will have to be remeasured from the start. After the completion of view 3 no erasures are possible as the event is then terminated.

Whenever any erasure is made the programme automatically supplements the information on the lamp display, which calls for the next measurement required, by printing the corresponding label. This is just the same message as is printed when the CHECK switch is used on its own.

(k) ACCEPT

The two remaining switches are less frequently used, but they have nevertheless a necessary function to fulfil. They enable the operator to omit certain measurements entirely, or to force acceptance of a track which has failed the curve fit test. When a track fails to be accepted when the END TRACK is given, then the operator can force its acceptance by using the ACCEPT switch. Its use is discouraged by recording a special message on the printer each time it is employed. It may also be used to bypass the measurement of a track. This could be necessary when the track is so confused with other tracks that it is unmeasurable. In this case no attempt is made to measure it, and the END TRACK is used with no points measured. The track is failed by the computer as is any track with

less than five points measured, and the operator forces the programme to move on to the next measurement by following the END TRACK with the ACCEPT setting.

(1) END VIEW

If an event is unmeasurable in one view, as may occur if the film is damaged or fogged, then the event can be measured in only two views. The measurement of the bad view can be bypassed by using the END VIEW switch. This may also be used part way through the measurement of a view, in which case the programme moves on to the next view, or in the case of view 3, terminates the event.

If more serious errors are made than can be recovered by any of these methods, then the event should be remeasured completely. This is achieved by starting the event again by raising the ACCEPT DATA interrupt. The incomplete event will be abandoned and the remeasurement initiated. During later processing the data for the incomplete event will be discarded. In practice however, the only time when this is done is right at the start of an event when an incorrect event number or topology coding has been given. It is rarely required at any other time as any mistakes that are made can be corrected by the appropriate choice of latching switches.

3.5 Preparation of data for further processing

The data from the on-line measuring system are written onto disk in essentially the same order as the measurements are made. For each event, heading information, coordinate measurements, and finally referencing information are output onto disk in card image form. The data are in

packed coding so that only 40 words are required for each 80 column card image. A maximum of 72 card columns are used for data, the last 8 columns being available for identification. Column 80 is used to tag the card with the number of the machine to which it refers.

The start of each event is indicated by a card image containing 444444 in columns 1 to 6. This is followed by a heading card containing all the usual heading information. As the event is measured, the coordinates are arranged on card images which are transferred to the disk as they are filled. Each card image accommodates six coordinate pairs. The end of the coordinates of the event is indicated by a card containing SSSSSS in columns 1 to 6, followed by several cards containing the referencing information for the event. This is in the form of a location table which enables the measurements of each point and track, in each of the three views, to be located within the bank of coordinates for that event. Finally, a single card containing ZZZZZZ in columns 1 to 6 marks the end of the data for that event.

During measuring several machines are usually being operated simultaneously. Data from each machine are output onto disk as they become available. Thus the card images containing the data from each machine are interspersed with those from any other machines that happen to be operating at the same time.

Later, when the machines are no longer being operated, the data are prepared for input to NUMAC. This is done by the programme SORT which selects card images by machine number, and then rearranges the cards for

each event into the order required for input to THRESH. This involves moving the referencing information so that it comes immediately after the heading card, and before all the coordinate measurements. SORT also discards incomplete events, and checks each event for errors which would make it unacceptable to THRESH.

The output from SORT is usually onto another disk, rather than back onto the same one. This is made possible by having two disk drives on the one system. The output disk is then taken to the terminal computer from where the data are input to NUMAC.

CHAPTER 4

PROGRAMMES FOR THE ON-LINE SYSTEM

No description of the on-line measuring system would be complete without an account of the main programming considerations. This chapter is therefore devoted to a description of the programmes which are used to control the system.

After the introduction in Section 4.1 a general discussion of interrupts is given in Section 4.2 as some knowledge of interrupt philosophy is considered to be desirable before attempting to understand the logic of the on-line system. The chapter then goes on to describe the logic of the programmes in detail, and gives a full account of how the transfer of data through the interface is controlled by the programme. Section 4.6 describes the curve fitting routine which is used to check the coordinate measurements of each track, and the way in which the data are output onto disk as described in Section 4.7.

4.1 Introduction

Before describing the programmes of the on-line system in any detail it will be helpful to discuss very briefly the overall logic of the system. Some idea of this has already been given in the previous chapter where the system was described from the operator's point of view. In this chapter however, the emphasis is on the computer and the programmes which monitor the operation of the measuring machines.

Chapter 3 also gave a full description of the hardware of the system, but a brief review is given here to draw attention to the main points

of interest to the programmer.

4.1.1 A review of the hardware

The measuring machines are connected to the IBM 1130 computer through the Westinghouse interface and an OEM channel. The CPU has 8K words of core and a cycle time of 3.6 μ sec. From the start it was appreciated that with this small size of core it would be necessary to make the most efficient use possible of the space available. In fact it soon became apparent that the only way in which this could be done was to write nearly all of the routines in Assembler language rather than in Fortran.

Each machine has a control panel through which the operator may communicate with the computer, and vice versa. Transfer of data into the computer is in the form of 16-bit binary words, of which there are seven associated with each machine. These correspond to the fifteen decade edge switches and ten latching switches on the control panel, and the X and Y coordinates.

Output to each measuring machine is in the form of a single 16-bit word. Fifteen bits are used to illuminate a pattern of lights on the lamp display of the control panel. The last bit may be used to sound the alarm.

The data from the machines are output onto magnetic disk. Two disk drives are available, but the second drive is required to take a separate disk for the output of data from the core store of the cosmic ray experiment.

Further information for the operators than can be shown on the lamp

displays, may be printed on the console printer. These messages are relayed by closed circuit television to individual monitors placed close to each measuring machine.

4.1.2 Logic of operation

The design of the on-line system was based upon two fundamental requirements. Firstly the system should check the data from the measuring machines as they are prepared, and secondly, the operators should be relieved of as much of the routine part of their work as possible. These requirements should not be realised at the expense of flexibility, nor should the operators be deprived of the facility of erasing or remeasuring points or tracks as this was found to be useful in the paper tape system.

It was decided that the computer, rather than the operators, would be in control of the measuring. This means that at the start of each event the operator must enter a coded topology number which enables the programme to determine exactly what measurements will be required for that event in each view. The programme then calls for each of the measurements in turn.

This method requires that the operators all conform to the same order of measurements for each event. The loss of freedom in this respect is not considered to be a disadvantage as little use was ever made of this freedom in the paper tape system. On the other hand, there are in fact two considerable advantages of having the computer in control. Firstly the operators no longer have to precede each measurement with an identifying label so that errors or omissions of labels can no longer occur, and secondly there is no

longer any danger of measurements being omitted.

The computer must then keep a list of all the measurements that are required for each of the possible six events currently being measured. It must also keep account of each measurement as it is made, and know how much of each event has been measured at any time.

4.1.3 Communication by interrupt

The measuring machines are not able, of themselves, to transfer data into the computer. This can only be done in response to certain programme instructions. Therefore, whenever an operator wishes to transfer data into the computer she presses one of her interrupt buttons to request the programme to read the data that are available. In response to the interrupt the programme gives the necessary instructions to read the seven data words from her machine.

The subsequent action of the programme depends upon what was requested by the operator. This is indicated by her choice of interrupt button (whether ACCEPT DATA or EXECUTE), and by the settings of the latching switches on her control panel. The programme finally informs the operator what she should do next by illuminating the appropriate lamps on her control panel.

4.1.4 Timing considerations

Since each operator must wait until her interrupt request has been attended to before proceeding with the next measurement, it is important that the computer's rate of response to each interrupt should be more than

fast enough to keep pace with the operator's rate of measuring. Most of the operators can work at the rate of about three interrupts in two seconds when they are measuring the 8 or 9 points which are required for each track. At other times the measuring rate is rather less than this. If the programme is to be able to respond to as many as six machines being operated simultaneously without causing the operators any noticeable delay, then the time required to service each interrupt should not be more than 0.1 to 0.2 seconds. Later in this chapter it will be seen how this restriction affects the programme.

4.2 Interrupts

The different instructions that make up any programme are normally executed sequentially. Any deviation from this requires the execution of a branch instruction which causes control to be transferred to an instruction other than the one immediately following.

However, there are situations in which it is necessary to interrupt the usual sequence of execution in order to allow the CPU to execute a different sequence of instructions of higher priority. When the CPU has completed the higher priority routine it usually returns to the original sequence of instructions and resumes execution from the point at which the interrupt occurred.

The use of such interrupts is fundamental to the operation of most computing systems, as it not only allows the operator to intervene in the execution of a programme, but it also makes it possible for most

input and output (I/O) devices to operate while the CPU continues the execution of a main programme. When the I/O device requires attention from the CPU it raises an interrupt which forces the CPU to give priority to the requirements of the I/O device.

Although interrupts are widely used by most computing systems, the present discussion refers specifically to the IBM 1130 system, and to the operation of the OEM channel in particular.

4.2.1 The nature of an interrupt

The actual operation of all I/O devices is controlled by the programme in the CPU. They cannot operate independently of the CPU, but rely upon it to control every function by means of control commands.

Usually the operation of an I/O device is initiated by the execution of a control command in the programme, and while the operation is being performed the main programme continues execution. When the required function has been completed, the I/O device informs the CPU by raising an interrupt.

The interrupt forces a branch out of the main programme into an I/O routine. The I/O routine then determines the cause of the interrupt and executes the next control command if this is required. After the interrupt has been attended to, a branch instruction in the I/O routine returns control to the main programme that was interrupted and its execution is resumed.

4.2.2 Interrupt levels

There are six different priorities or interrupt levels within the 1130 system, and each I/O device is assigned to one or more of these levels. In Table 4.1 the interrupt levels used by some of the more common I/O devices are shown.

Table 4.1

Table of I/O devices and interrupt levels of IBM 1130 system

Interrupt level	Device
0	1442 Card Read Punch (column interrupt)
1	1132 Printer
2	* Disk Storage * Storage Access Channel (SAC) (for second disk drive)
3	1627 Plotter * OEM Channel SAC
4	1442 Card Read Punch (operation complete interrupt) * Keyboard/Console Printer * 1134 Paper Tape Reader * 1055 Paper Tape Punch 2501 Card Reader 1403 Printer 1231 Optical Mark Page Reader SAC
5	* Console (INT RUN and PROGRAM STOP) SAC

* Items marked with an asterisk are part of the on-line system.

Those devices using the interrupt levels of high priority are able to interrupt the servicing of low priority interrupts in the same way that an interrupt on any level can interrupt the execution of a main programme. In this respect level 0 has priority over all the others, while level 5 has the least priority.

Associated with each interrupt level there is an interrupt level subroutine. The six core locations 8 to 13 are reserved for the starting addresses of the six interrupt level subroutines, so that when an interrupt occurs the CPU can force a branch to the correct subroutine.

4.2.3 Interrupt servicing

When an interrupt occurs, the CPU forced branch is not executed until the main programme instruction currently being executed is completed. The branch, which is in the form of a BSI instruction (branch and store instruction address register), then stores the address of the next main programme instruction in the first word of the interrupt level subroutine, and transfers control to the second word. The storing of this address enables the subroutine to return control to the correct instruction in the main programme when it has finished servicing the interrupt.

Before the I/O routine attends to the interrupt it must save the contents of all of the machine registers that it will use during the servicing of the interrupt. These values are stored in core, and are restored to the registers before control is returned to the main programme. If this were not done it would not be possible for the execution of the

main programme to be resumed as vital information would have been destroyed.

Associated with each of the interrupt levels 1 to 5 there is an interrupt level status word (ILSW). Each ILSW is a 16-bit register which may be read into the accumulator in the CPU. The different bits in each ILSW are assigned to the different I/O devices that can interrupt on that level (see Figure 4.1).

Similarly there is a device status word (DSW) associated with each I/O device. Each DSW has at the most 16 bits, which may be read into the accumulator, or read and reset by the programme. The different bits in the DSW are used to indicate different conditions of the device.

The determination of the cause of an interrupt usually involves the examination of both the ILSW and the appropriate DSW. The ILSW is examined first to determine which device raised the interrupt. Then, when the device is known, the appropriate DSW is examined to find out what condition of the device caused the interrupt to be raised.

The programme can then respond to the interrupt by executing the required control command. For example, the interrupt might have been given during the transfer of data into the computer. As each data word becomes available for reading in, an interrupt is raised. The programme should respond by executing a read instruction. Similarly, during the

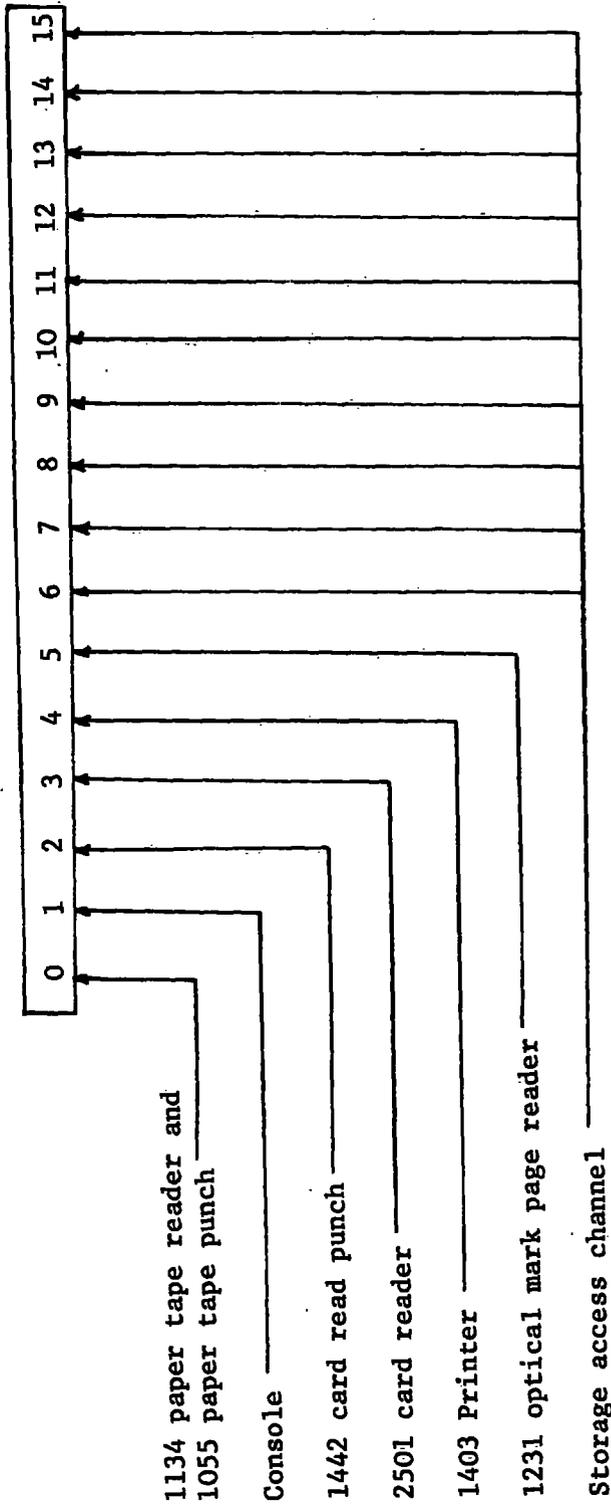


Figure 4.1 Interrupt level status word for level 4

output of data, the device will interrupt when it has received each data word to indicate that it is ready to receive the next word. In this case the programme should respond by executing the next write instruction.

4.2.4 OEM channel

In the case of the OEM channel the interrupts are raised on level 3. As can be seen from Table 4.1 two other devices, namely the storage access channel and the 1627 plotter, may use level 3 interrupts. However, the Durham on-line system has neither of these devices, on this level, so that the OEM channel is the only device in the system that can interrupt on level 3. It is therefore not necessary to examine the level 3 ILSW since any interrupt on that level must have come from the OEM channel.

There are effectively two device status words associated with the OEM channel, although in the Durham system only one of these is used. It is more often referred to as the interrupt register and is described in some detail in Section 4.5.1.

4.2.5 Inhibition of interrupts

During the servicing of any interrupt, further interrupts on the same, or lower, priority level are inhibited. This is because once the interrupt is raised, it remains on until it is switched off by the execution of a BOSC instruction (branch or skip on condition), and it is not until the interrupt has been switched off that another interrupt on the same, or lower, priority level will be recognised. Usually the

BOSC instruction to turn off the interrupt is executed when it is required to return control to the main programme after the interrupt has been serviced. In this way the programme can complete the servicing of one interrupt before attending to the next one.

It has already been explained how the use of interrupts makes it possible for the main programme to continue execution while an I/O device is operating. However, it is also possible for an I/O device to operate simultaneously with the execution of an interrupt subroutine, provided the I/O device uses interrupts of a higher priority level. If this were not the case, then the interrupts from the I/O device would not be recognised and so the programme would not be able to give the necessary control commands when required. For example, a level 3 routine may be executed simultaneously with the operation of disk input or output which uses level 2 interrupts, but not with the operation of the console printer which uses level 4.

4.2.6 Operator initiated interrupts

Although the most frequent use of interrupts is to inform the CPU when a CPU initiated operation has been completed, interrupts may also be used to allow operator intervention in the execution of a programme. One example of this is the INT REQ (interrupt request) key on the console keyboard of the computer. When this is pressed, a level 4 interrupt is raised and the usual CPU forced BSI instruction is executed. The subsequent action of the CPU depends on the programme itself and so may be determined by

the user.

Interrupts of this kind may be referred to as operator initiated interrupts as they are the direct result of an action by the operator. In the on-line system interrupts of this kind play a very important part.

There are two interrupt buttons on each measuring machine, and since these are connected to the CPU through the OEM channel, the interrupt that is raised is on level 3. The operator uses these interrupts to inform the computer when she has data available for reading in, and so the computer must be programmed to respond to the interrupt by reading in the data from her machine.

4.3 Interrupt logic

In the on-line system the computer is used not just for the collection of data as they become available, but also as a monitor which both controls the measuring machines and checks the data as they are prepared. Basically the background programme is in a WAIT, pending an interrupt from one of the measuring machines.

The response to the interrupt involves the reading in of the seven data words from the machine that raised the interrupt. The subsequent action of the programme depends upon which of the two interrupt buttons was pressed, and which of the latching switches on the control panel were set. In each case the final response to the interrupt involves the output of a digital output word (DOPW) to the lamp display on the control panel of the machine. This completes the interrupt servicing, and a BOSC

instruction is given to turn off the interrupt.

4.3.1 Response to an interrupt

The pressing of an interrupt button on any of the machines causes a programme interrupt on level 3. By hardware control a forced BSI instruction is executed which transfers control to the routine OEMI, which services the OEM channel interrupts.

The first response to any interrupt is always to save the machine registers, and then to determine the cause of the interrupt. In the on-line system there is never any doubt about which device can have raised an interrupt on level 3, as the OEM channel is the only device in the system which uses that level.

The programme does however, have to determine whether it was raised by one of the measuring machines, or by the core store of the cosmic ray experiment. This is done by examining the interrupt register which is one of two device status words associated with the OEM channel. (It is described in detail in Section 4.5.1.)

In the case of an interrupt from one of the measuring machines it is necessary to determine which machine was requesting attention. A special register within the interface, the Westinghouse device status word, holds this information. By examining this word, the programme can determine not only which machine raised the interrupt, but which interrupt button was used (see Section 4.5.2).

Once this is known, the programme can read in the seven data words

from that machine. These are stored in the high core storage addresses where they may be accessed by any of the several routines which may be called upon to process them.

If the ACCEPT DATA button had been used to raise the interrupt the sign bit (bit 0) is set in the word which holds the settings of the latching switches. Another part of the programme later examines this word to determine what was requested by the operator, and to select the appropriate routine to process the data. This word, which is referred to as the function request word, is illustrated in Figure 4.2.

4.3.2 Selection of data processing routine

The selection of the routine to process the data is illustrated in Figure 4.3. It involves the examination of the function request word.

The programme begins by examining the sign bit to determine which interrupt button had been used. If the sign bit is set the interrupt was raised by the ACCEPT DATA button, in which case the request is for the start of a new event and the routine NEVT is called.

If the EXECUTE interrupt button had been used the sign bit is not set and the programme then scans bits 6 to 15 of the function request word to determine which latching switches had been set. A series of SLA (shift left accumulator) instructions is used for this. Except for two special cases only the left-most bit that is set is noticed by the programme. The special cases (see Figure 4.3) are the ERASE TRACK and the CHECK settings.

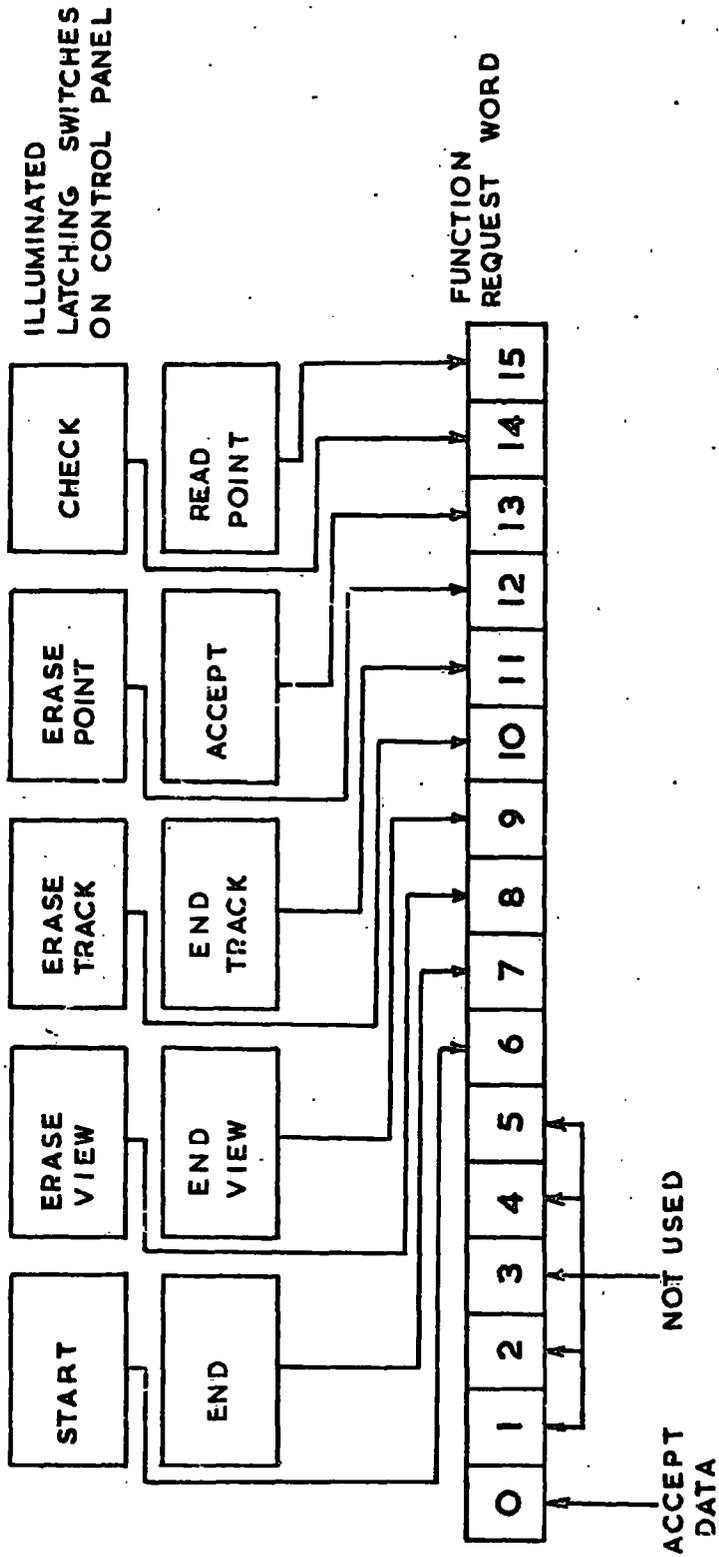


FIGURE 4.2 — FUNCTION REQUEST WORD —

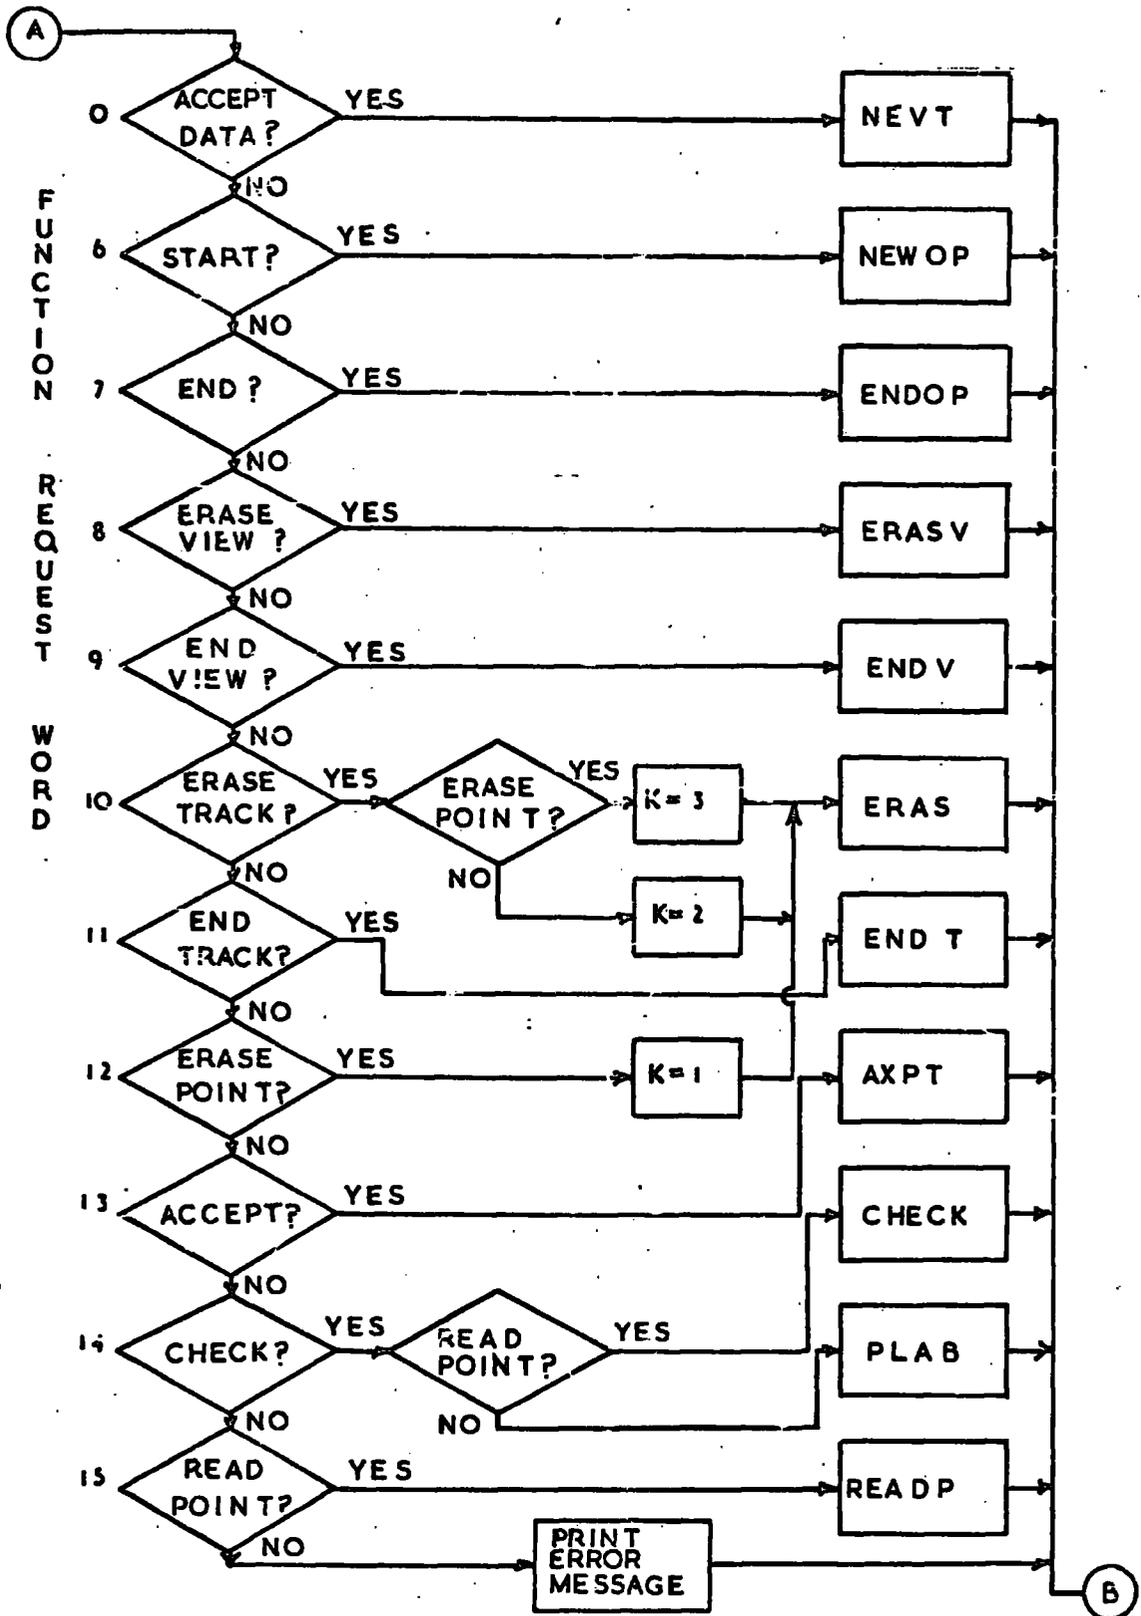


FIGURE 4.3 - SELECTION OF DATA PROCESSING ROUTINE -

As the EXECUTE interrupt is inhibited when none of the latching switches is on, it might be considered redundant to check for the final possibility, READ POINT, when all the others are off. However, it is possible, though unlikely, for an operator to raise an interrupt with her EXECUTE button, and then reset whichever latching switch was set on her control panel before the programme reads these settings. In such a case, it will appear to the programme as if none of the switches had been set. Instead of assuming that the READ POINT switch was set, the programme causes an error message to be printed to inform the operator that her last request was not understood.

The sequence of instructions which is responsible for examining the function request word and for calling the appropriate routines to process the data, is part of a controlling routine called MAINE. The other functions carried out by the rest of this routine are described in Section 4.3.5.

4.3.3 Data processing routines

Any one of eleven routines may be selected by MAINE to process the data which have been read in from a machine in response to an interrupt. The names of these routines all appear in Figure 4.3. This logic diagram also indicates the type of request that each routine has to handle.

Although the main function of each routine is peculiar to itself, there are some minor functions which are common to all of them. The first responsibility of each routine is to check that the request indicated is a valid one. For example the END TRACK setting is not valid when the

operator is being asked to measure an individually labelled point, or again the ERASE POINT is not valid just after a track has been completed and accepted.

If the request is not valid the routine initiates the printing of a coded error message which informs the operator of the nature of the fault.

If the request is valid, the routine will proceed to fulfil its main function which may involve the setting up of a new event, or the storing of a pair of coordinates, or the checking of the coordinates of a completed track, or the erasure of the last point measured, and so on.

In every case, the routine has to prepare the digital output word which will be used to illuminate the lamp display at the conclusion of the interrupt servicing. This word will include the alarm bit if the routine has detected any error condition such as an invalid request, or if the data which have been examined are found to be unacceptable as in the case of a track which fails the curve fitting test.

The routine finally returns control to MAINE as shown in Figure 4.3. MAINE then transfers control to OEMI as this routine includes the sequence of instructions which control the output of the DOPW to the lamp display.

4.3.4 The OEM channel routine

The routine which services the level 3 interrupts from the OEM channel is called OEMI. This routine is responsible for controlling all movement of data through the interface. It saves the machine registers when an interrupt occurs, determines which machine raised the interrupt, and reads

in the seven data words from that machine. After the data have been processed by MAINE, OEMI routes the DOPW to the lamp display of the machine and turns off the interrupt.

(a) Handling of different interrupt levels

After an interrupt has been raised it remains on until it is turned off by the execution of a BOSC instruction. So long as it remains on, any further interrupts on the same, or lower, priority level will be inhibited. Thus, if the routine OEMI does not execute a BOSC instruction until the servicing of each interrupt has been completed, there will be no possibility of one machine interrupting the servicing of another machine.

The situation is complicated however, by allowing the response to certain interrupts to require the printing of a message on the console printer. This may be an error message which is given when an invalid interrupt is made, or when the coordinates themselves are not acceptable. A printed message is also given when the operator requests more information than is given on her lamp display about what she should do next. Another use of the printer is to log the start of each event.

In each of these cases the use of the printer is required as part of the response to an interrupt. Now the console printer uses level 4 interrupts and so cannot operate simultaneously with the execution of a level 3 routine as its interrupts would not be recognised. It is therefore necessary for a BOSC instruction to be given during the servicing of the level 3 interrupt to turn off that interrupt level before the instructions are given to initiate the operation of the console printer. The servicing

of the level 3 interrupt then continues at main level although control is not returned to the background programme that was interrupted until the servicing is completed.

Once the BOSC instruction has been executed it immediately becomes possible for another machine to raise an interrupt. In order to avoid the confusion that would arise if the programme were forced to abandon the servicing of one interrupt to attend to another interrupt, a 'busy flag' is set by the interrupt programme before the BOSC instruction is given, and it is not cleared until the servicing of the interrupt has been completed. Any interrupts occurring while the 'busy flag' is set are noted, but are not attended to until the servicing of the current interrupt has been completed.

Before completing the servicing of each interrupt the programme returns to interrupt level 3 for the output of the DOPW to lamp display. The only way of switching on an interrupt level is to raise an interrupt, and in this case a programme initiated interrupt, or forced interrupt, is required. The way in which this is done is described in Section 4.5.1. The DOPW is then output to the lamp display, and the 'busy flag' is cleared.

Provided no further interrupts have occurred in the meantime, the final BOSC instruction is given to switch off the interrupt level again and return control to the main programme. If however, further interrupts have occurred, the BOSC instruction is not given immediately. The interrupt level remains on and the programme proceeds to service the second interrupt as if it had just occurred.

A simplified form of the logic of the OEMI routine is shown in Figure 4.4 (the complete flow diagram is shown in Appendix A), and that part of the servicing which takes place at main programme level is illustrated in Figure 4.5. An examination of Figure 4.5 reveals that the 'busy flag' is also used to inform this part of the programme that it is required to process the data which have just been read in. This diagram is also simplified for the sake of clarity but is shown in its complete form in Figure 4.6.

(b) Cycling of priorities

So long as each interrupt from a measuring machine is serviced completely before any other interrupt is raised, there can be no question of one machine receiving more rapid attention than any other. However, when several machines are being operated simultaneously it frequently happens that interrupts are raised while the computer is already busy servicing another interrupt.

When the servicing of the first interrupt is completed, the programme examines the Westinghouse device status word (see Section 4.5.2 and Figure 4.10) to see if any other interrupts have been raised in the meantime. As this word (the WHDSW) is examined by shifting its bit pattern to the left until a "1 - bit" is shifted into the sign position (bit 0) then the two bits for machine 6 will be detected before any of the bits for the other five machines.

If the WHDSW is read and examined in this way after each interrupt has been serviced, machine 6 will always have priority over all the other

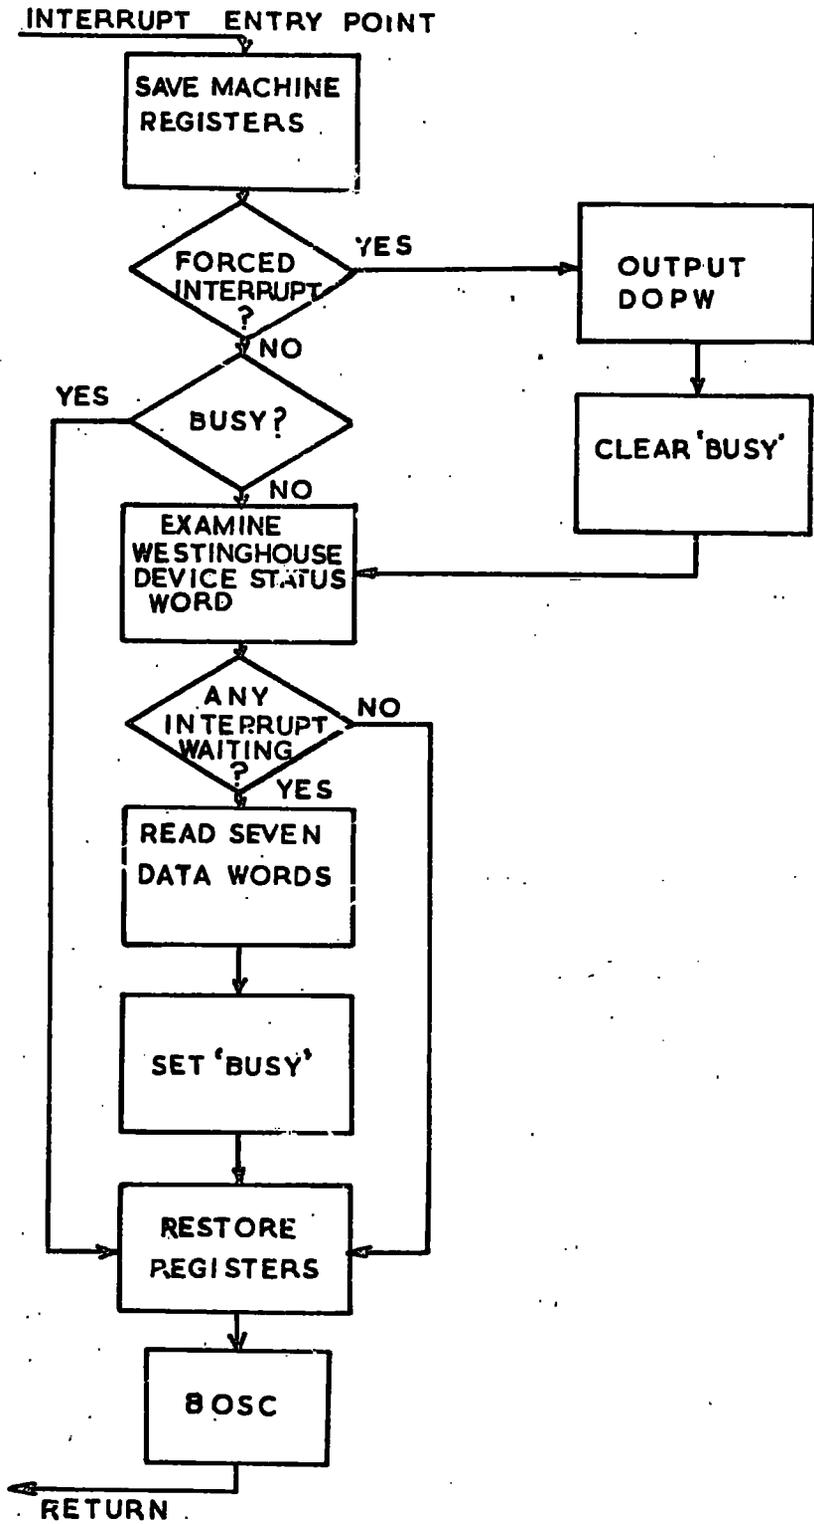


FIGURE 4.4 — OEMI ROUTINE (SIMPLIFIED) —

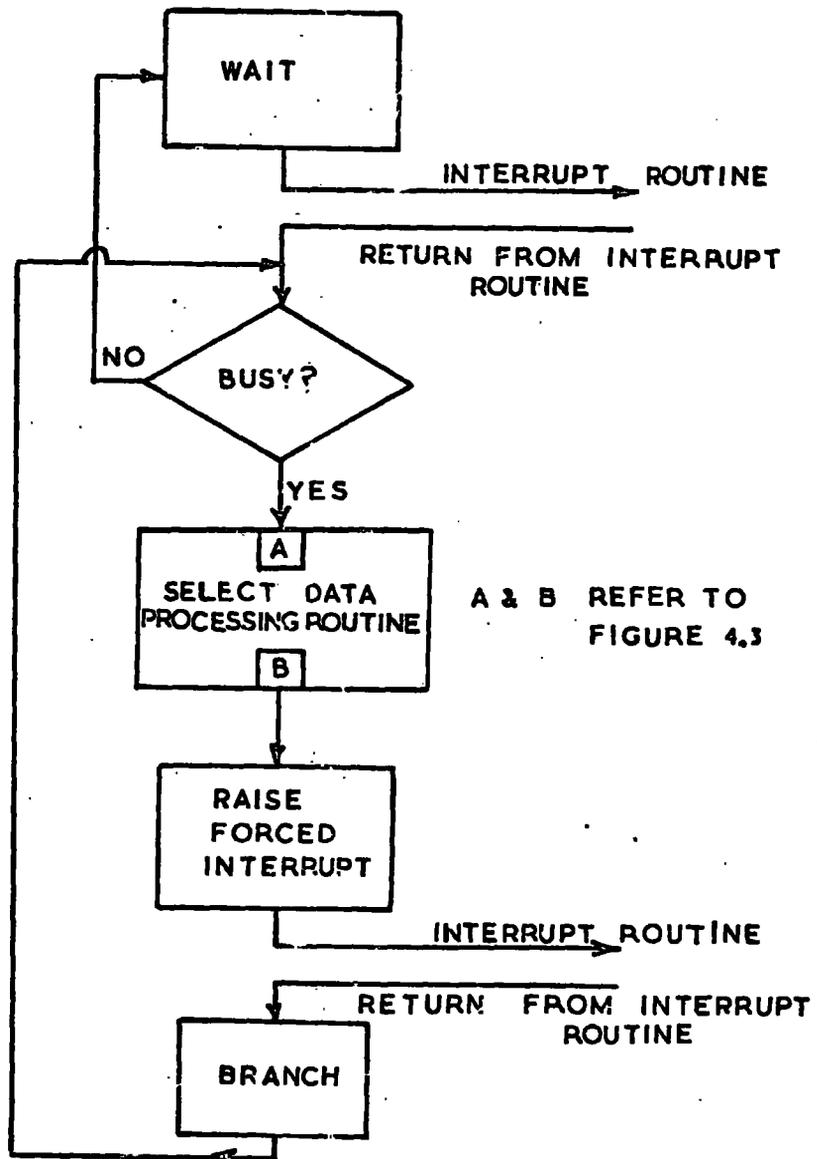


FIGURE 4.5 - **MAINE (SIMPLIFIED)** -

machines, and it could happen that an interrupt from machine 1 could be left unattended for some time.

In order to reduce this lack of balance in priorities, the WHDSW is not read again immediately following the completion of the servicing of each interrupt. After the WHDSW has been examined and the left most bit has been detected, the remaining bits of the WHDSW are preserved for future examination. Then, when the servicing of the first interrupt has been completed, the programme examines the remaining bits of the WHDSW that were preserved after the previous examination. In this way the interrupts corresponding to any other bits that were set in the WHDSW when it was last read in, will be serviced before the WHDSW is read in again. While still giving some priority to machine 6, this method does ensure that no machine can be ignored indefinitely.

(c) Contact bounce of interrupt buttons

During the initial testing of the on-line system it was found that the interrupt buttons were occasionally raising two interrupts at a time. Special switches were being used which were designed to prevent contact bounce, but no amount of careful handling could eliminate this effect altogether.

In the case of the ACCEPT DATA interrupt button which is used to indicate the start of a new event, a double interrupt creates no problem as it only causes the programme to initialise the event twice. Two sets of the heading information for the event are prepared for output onto disk, but the first is discarded when the data are edited by SORT before being input to NUMAC.

A double interrupt from the EXECUTE button, on the other hand, is most undesirable as two responses to the same interrupt would interfere with the correct measurement of the event. Depending upon which latching switches on the control panel were set at the time, this might or might not have resulted in the printing of an error message, and so the operator might or might not have been aware that a double interrupt had occurred.

Unlike the ACCEPT DATA button, the EXECUTE button not only raises an interrupt, but also causes the contents of the X and Y coordinate registers to be transferred into two buffers. This is to allow the operator to move away from the point on which she has just set before the coordinate measurements are actually read by the computer.

Whenever a double interrupt from the EXECUTE button occurs, the coordinate buffers are only loaded with the coordinates on the first interrupt. This means that when the coordinate buffers are read by the programme in response to the second interrupt both coordinates are found to be zero. The computer is therefore programmed to examine the coordinates that are read in after each interrupt, so that if both are found to be zero the interrupt is not serviced in the usual way but is ignored, except for the printing of a message which simply records that a double interrupt has been detected. The alarm is not sounded and the operator is able to proceed as if nothing had happened.

4.3.5 The background programme

The main programme which is in core during on-line control of the

measuring system is called MAINE. Some parts of its function have already been described, in particular the part which selects the appropriate routine to process the data from each interrupt.

The programme MAINE does however, have more to do than is indicated in Figure 4.5, where it appears as if the background programme is nothing but a WAIT, followed by an examination of the 'busy flag'.

The complete background programme is illustrated in Figure 4.6. Here it can be seen that, before checking the 'busy flag', MAINE reads the console entry switches (CES) whenever it leaves the WAIT. The CES are a set of 16 toggle switches on the computer console, which are numbered 0 to 15 and whose settings may be read by the programme. Each reading is compared with the previous reading which is recorded in one of the high core storage addresses.

Switch 15 is examined first to see if its setting has changed since the previous time the CES were read. This switch is used to indicate whether or not the interrupts from the core store of the cosmic ray experiment are to be serviced. If it is not on, the programme will not read in any data from the core store if it raises an interrupt, but will only clear the interrupt. When switch 15 is switched on MAINE calls a routine called CRINT to perform certain initialising functions that are required before data from the core store can be handled correctly. Similarly there is a shutting down procedure called CRSD which is called by MAINE when switch 15 has been turned off.

In much the same way switches 1 to 6 are used to indicate which measuring machines are currently being operated. The programme checks for

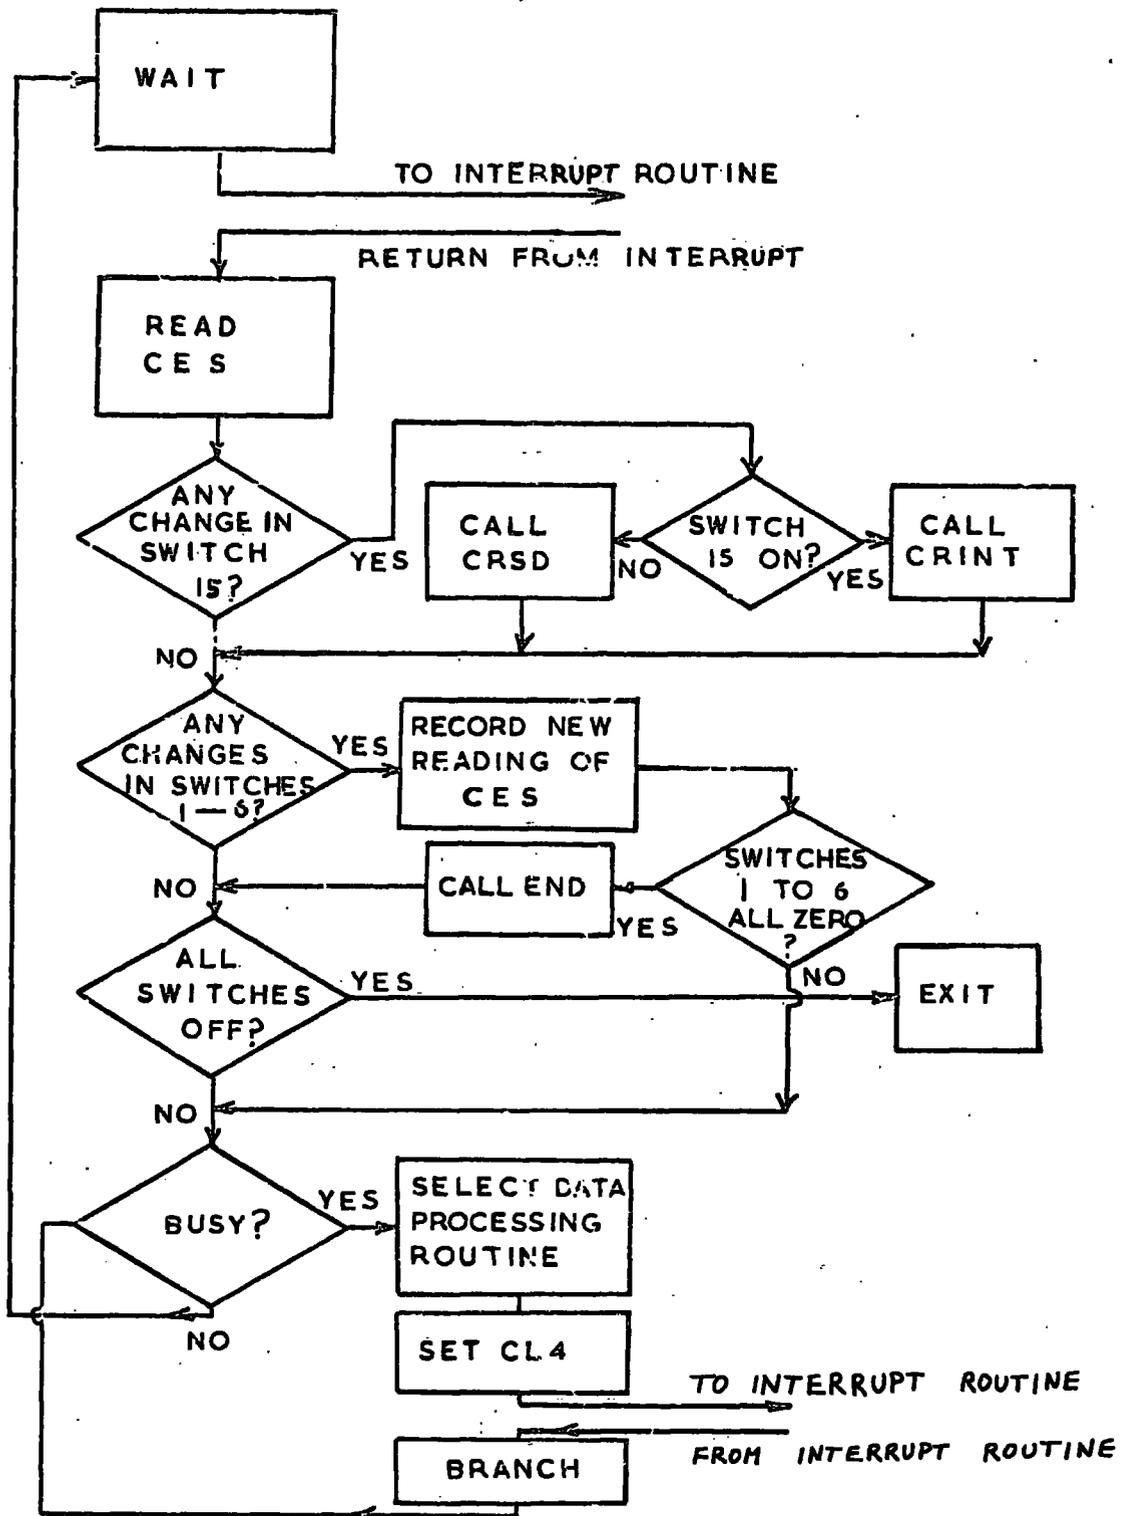


FIGURE 4.6 — MAINE (COMPLETE) —

any change in the settings of those switches. If there is a change the new settings are recorded. Any interrupts from a machine whose corresponding console entry switch is not set, will not be serviced. The data words will still be read in as usual so that the interrupt will be cleared, but no examination of the data will take place and no DOPW will be output.

If at any time it is found that switches 1 to 6 have all been switched off, then the programme calls a closing down routine called END. Once END has been executed, no more data from the measuring machines can be accepted. If it is required to resume measuring, then the programme must be reloaded.

If switch 15 is also off when switches 1 to 6 are turned off, then neither the measuring machine nor the core store interrupts are to be serviced and so the programme is terminated by the execution of an EXIT instruction.

If, however, some of the switches are on, the programme checks the 'busy flag' to determine whether or not there are data from a measuring machine which are waiting to be examined. If the 'busy flag' is not set, the programme returns to the WAIT until the next interrupt occurs, or until the PROGRAM START button on the computer console is pressed.

If the 'busy flag' is set, MAINE proceeds as described in Section 4.3.2 to examine the function request word to determine which data processing routine to call. This data processing routine prepares a DOPW before returning control to MAINE, which then transfers control to OEMI for the output of the DOPW. After the DOPW has been routed to the appropriate lamp display and the 'busy flag' cleared, OEMI checks to see if any more

interrupts have occurred in the meantime. If there have been no further interrupts OEMI returns control to MAINE which then returns to the WAIT.

On the other hand, if another interrupt has occurred in the meantime, OEMI turns its attention to this second interrupt as soon as it has finished with the first. It reads in the seven data words from the second machine that interrupted, and sets the 'busy flag' again, before returning control to MAINE. MAINE detects the 'busy flag' as it did for the first interrupt, and once again it examines the function request word to determine which routine to call to process the data.

This cycle continues until all outstanding interrupts have been serviced. When this happens OEMI clears the 'busy flag' before returning control to MAINE after completing the servicing of the final interrupt. Only then does MAINE return to the WAIT.

4.4 Organisation of data in core

In order to keep a record of the progress made throughout the measurement of each event, the programme holds certain details of the event in core. Some of the data for the event have also to be held in core. In particular the coordinates of each track have to be held until the track is completed so that they may be checked by the curve fitting routine before being output onto the disk. Certain referencing information is also required to enable the coordinates to be associated later with the correct points and tracks.

Most of these data are held in a common area at the top of core where they are accessible to all the various routines which make up the controlling programme for the on-line system. This makes it possible to keep the passing of arguments between the different routines down to a minimum, and so results in a valuable saving of core space.

4.4.1 Common

Of the 8192 words which make up the core of the CPU, just over 2000 words are allocated to "common". Those words, which are located at the high core storage addresses, contain essentially all the data that are held in core during the measurement of each event.

Most of common is in the form of arrays with space allocated to each of the possible six measuring machines. It also includes two 320 word disk output buffers, one for the output of data from the measuring machines, and one for the output of data from the core store of the cosmic ray experiment.

A list of the different variables in common, and their dimensions is given in Table 4.2. The addresses are shown in their hexadecimal form (base 16), and are given relative to a core of 32K words. These addresses are also applicable to a core of 8K words because of the 'wraparound' method of interpreting core storage addresses. The reason for giving the addresses in this form is to make the programmes compatible with a 16K or a 32K machine, so that, should the CPU ever be updated by the addition of extra core, the programmes for the on-line system could continue to be used without modification.

Table 4.2
List of variables in common

LABEL	No. OF WORDS	CORE ADDRESSES	DESCRIPTION
DATAT (8)	8	7FFF - 7FF8	Data Table
DOPW	1	7FF7	Digital output word
FRW	1	7FF6	Function request word
CRES (327)	327	7FF5 - 7EAF	Cosmic ray reservation
DTOP (43)	43	7EAE - 7E84	Digested topology
OAC (324)	324	7E83 - 7D40	Output area and counters
DATE (2)	2	7D3F - 7D3E	Date
CARDS (240)	240	7D3D - 7C4E	Cards
CARD (72)	72	7C4D - 7C06	Card
FDCLS (3,4,6)	72	7C05 - 7BBE	Fiducial references
PNTS (4,8,6)	192	7BED - 7AFE	Point references
TRCKS (7,10,6)	420	7AFD - 795A	Track references
PNO (6)	6	7959 - 7954	No. of points on current track
TRACK (2,15,6)	180	7953 - 78A0	Coordinates of current track
CPNO (6)	6	789F - 789A	No. of coordinate pairs
FDCL (2,4,6)	48	7899 - 786A	Fiducial coordinates
APEX (2,6)	12	7869 - 785E	Apex check coordinates
NPL (6)	6	785D - 7858	Number of point labels
NPL (6)	6	7857 - 7852	Number of track labels
POSTN (6)	6	7851 - 784C	Position
CKNPL (6)	6	784B - 7846	Number of point labels at last check
CKNTL (6)	6	7845 - 7840	Number of track labels at last check
CKPOS (6)	6	783F - 783A	Position at last check
VIEW (6)	6	7839 - 7834	View
OPER (6)	6	7833 - 782E	Operator and experiment number
ENDF (6)	6	782D - 7828	End flag
APXF (6)	6	7827 - 7822	Apex flag
STF (6)	6	7821 - 781C	Short track flag
FDCLF (6)	6	781B - 7816	Fiducial flag
LAB	1	7815	Label argument

Note Core addresses are shown in hexadecimal form

Starting from the highest addresses in core, the first 8 words are used for the number of the machine currently being, or most recently having been serviced, and the seven data words which were read in from that machine. The location /7FF7 is used to hold the digital output word which is prepared in response to each interrupt. The settings of the console entry switches, which are used to indicate which machines are being operated, are recorded in location /7FF6.

The next 327 words are reserved for the handling of data from the core store of the cosmic ray experiment. They include a 320 word disk output buffer.

During the measurement of each event the programme keeps a record of all the measurements required for that event. This information, for all of the six machines, is stored in the array DTOP which is described in the next section.

This is followed by the disk output buffer which is used for the data from the measuring machines. A few words are also reserved for some counters which are used in connection with the output of data onto disk.

Two words are required for the date which is stored in binary coded decimal.

The 240 word array called CARDS consists of six 40 word card images, one corresponding to each of the six measuring machines. Data are collected in these card images in packed coding, two characters per word, before being transferred to the disk buffer for output onto disk. The single 72 word array CARD is also used for preparing data for output. Full

details of how the data are prepared for output are given in Section 4.7.

During the measurement of each event the coordinates are transferred to the disk in essentially the same order as the measurements are made. In order to associate the coordinates later with the appropriate points and tracks, certain referencing information is collected in the arrays FDCLS, PNTS and TRCKS throughout the measurement of each event. The arrays PNTS and TRCKS also hold the labels of the points and tracks which are required for the events currently being measured. Sufficient space is allocated to these arrays to allow for a maximum of eight individually labelled points, and ten tracks to be measured for any event.

The array PNO keeps a record of the number of points already measured on the current track. As usual there are six words in this array, one corresponding to each of the six machines. The coordinates of each track are stored in the array TRACK until the measurement of the track has been completed. Acceptable measurements are then prepared for output onto disk, while unacceptable measurements will be overwritten as the track is remeasured. A maximum of 15 points may be measured on each track.

During the measurement of each event a record of the number of coordinate pairs that have been prepared for output is kept in the array CPNO.

The measurements of the four fiducials in each view are held in the array FDCL until the view is completed. This makes it possible for the relative positions of the fiducials to be checked before the complete view is accepted, although there is insufficient core space available at the moment to accommodate the checking routine.

The next array APEX, was included to hold the coordinates of the apex setting values from each of the six machines. This setting value is entered at the start of each view and any later measurements of the apex are compared with the setting value to determine whether or not the digitisers on the measuring machines are counting correctly. Later on it was decided to insist upon the same setting value for each of the measuring machines so that this array is no longer required. However, it has not been removed as it may be found desirable at some later stage to vary the setting value from one machine to another.

Throughout the measurement of each event the programme keeps a record of the progress made. The array POSTN records a number corresponding to the measurement currently being requested (see Section 4.4.3). NPL and NTL record the number of individually labelled points and tracks respectively which have already been accepted for the view currently being measured.

The arrays CKNPL, CKNTL and CKPOS preserve the values that were in NPL, NTL and POSTN respectively when the last successful apex check was made. This enables the programme to return to that position in the sequence of measuring when the apex check was last accepted, should a subsequent failure occur. The values of CKNPL, CKNTL and CKPOS are updated after a specific request for an apex check has been successful, or after the measurement of a track has been accepted by the curve fitting routine.

The number of the view currently being measured is recorded in the array VIEW, and OPER holds the current values of both the experiment number and the operator number which are read in from the decade edge switches

of the control panel when the operator signs on at the beginning of her shift.

The final four arrays are flags which are used by the programme for various purposes. For example, the use of ENDF provides a safeguard which prevents an operator from signing off by accident when she is in the process of measuring an event. STF is used to allow the operator to add more points to a track which has failed because insufficient points were measured or because the sagitta was too small. The use of STF also makes it possible for the programme to accept a track with no measured points.

The last word in common, called LAB, is used as a return argument by two routines whose task is to determine what kind of measurement was last requested, whether it was for a point or a track, or if it was for a fiducial mark and if so which fiducial.

There are a few other arrays which are not in common, but which are located within the subprogrammes that require them. Each is used only by the subprogramme of which it is a part, and so does not need to be incorporated in common.

4.4.2 The start of an event

Some of the arrays in common require to be initialised at the beginning of each event. In particular the array DTOP is set up to indicate which measurements are required for the event just being started. Then the labels corresponding to the points and tracks that are to be measured are set up in the arrays PNTS and TRCKS.

This is all done by the routine TOPOL which is called by NEVT after

an ACCEPT DATA interrupt has been raised. The arrays are initialised as TOPOL interprets the topology coding that was read from the second row of decade edge switches on the control panel. Of the seven digits available, only five are used at present, and these are interpreted as illustrated in Figure 4.7.

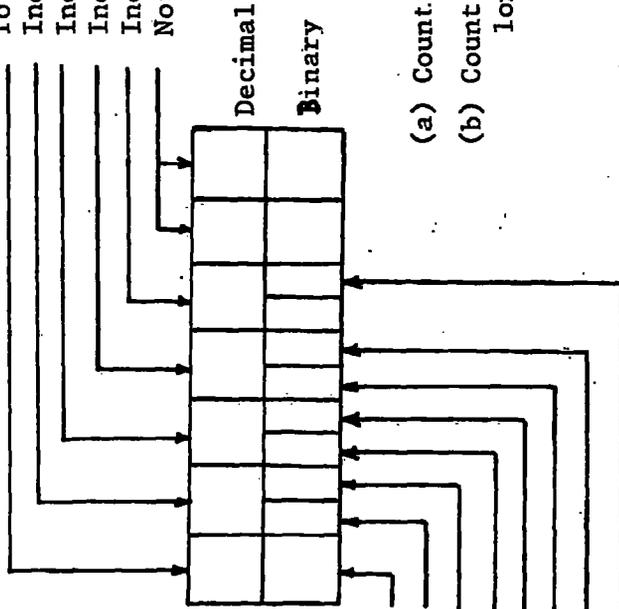
Some form of packing was required in order to code the whole event in seven digits. The method used, that of coding two distinct quantities on each of three of the switches, was chosen as it is particularly easy to decode. Each decade switch is read into the computer as four bits of a binary word. The four bits may be divided into two groups of two bits each. In the case of short stopping tracks the first two bits give the binary count of the number of short stopping protons, and the second two give the binary count of the number of short stopping pions. Since the maximum value that can be represented on a decade switch is 9, this method imposes some restriction on the possible combinations of special tracks for any particular event. In practice however, an event is rarely found which cannot be coded in this way.

Before setting up the list of measurements corresponding to the topology given, the topology is checked to see if it is acceptable. This involves making sure that the topology corresponds to an event which the on-line system can handle, and checking that the topology is self-compatible. Examples of unacceptable events would include any requiring more than two short stopping pions to be measured, or more than eight ordinary (not stopping) tracks. In short, any event which cannot be represented by some



Settings on edge switches

- Total number of secondary tracks
- Indicates number of short stopping tracks (a)
- Indicates number of long stopping tracks (a)
- Indicates number of kinks (b)
- Indicates number of V zeros
- Not used



Interpretation by programme

- Total number of secondary tracks
- Number of short stopping protons
- Number of short stopping pions
- Number of long stopping protons
- Number of long stopping pions
- Number of long kinks
- Number of short kinks
- Number of V zeros

- (a) Count 1 for each pion, 4 for each proton
- (b) Count 1 for each short kink, 4 for each long kink.

Figure 4.7 - Interpretation of topology coding -

combination of the 43 measurements listed in Table 4.3, cannot be handled by the on-line system. An example of an event which is not self-compatible would be one which requires the measurement of two short stopping tracks and one long stopping track, and yet was coded as having only two secondary tracks altogether.

Provided the topology is acceptable, TOPOL prepares a list of all the measurements that are required for the event. The array DTOP (see Section 4.4.1) which has one word corresponding to each of the 43 measurements that may be included in any event, is used. Only six bits of each word are required, one associated with each of the six measuring machines (see Figure 4.8). The appropriate bit is set in each word which corresponds to a measurement that is required. Some examples of the settings in DTOP that correspond to events of different topologies are shown in Table 4.3.

The measurements are requested in the order that they appear in the complete list of 43 measurements in Table 4.3, and the coordinates are output onto disk in essentially the same order. On the disk the measurements for each event appear in the form of a bank of some 100 or more coordinate pairs. To enable these measurements to be associated later with the correct points and tracks, the programme records the position, within this bank, of the coordinates of each measurement, as they are prepared for output. These addresses are stored in the arrays FDCLS, PNTS and TRCKS (see Section 4.4.1), which have sufficient space allocated to them to record this information for six different events from the six measuring machines. As each event is started TOPOL reinitialises the appropriate words in these arrays, and sets up in PNTS and TRCKS the labels corresponding to the measurements that are required.

Table 4.3

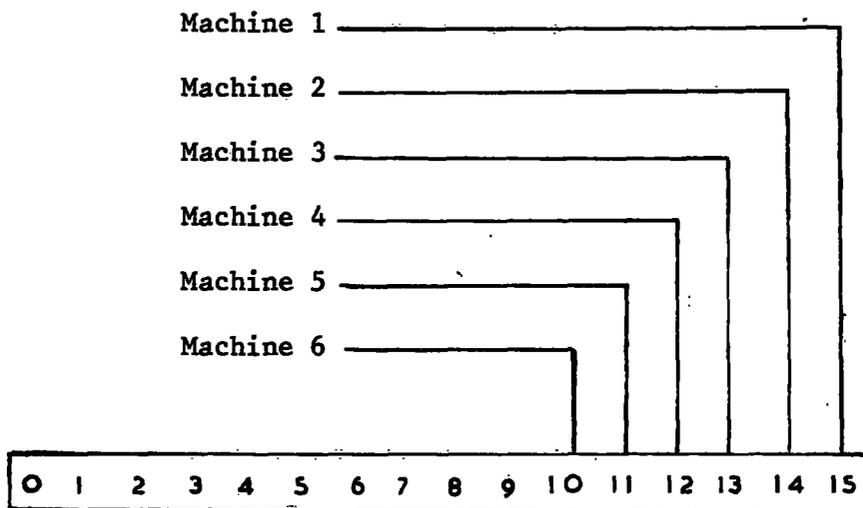
103

Examples of settings in the array DTOP corresponding to events of different topologies

M/C 1(bit 15)Top. 3040000 M/C2(bit 14)Top. 2400000 M/C3(bit 13)Top. 4010000

POSTN	LABEL	DTOP															DESCRIPTION	
		0	1	2	3	4	5	6	7	8	9	10	11	12	13	14		15
1	AA														1	1	1	Apex
2	A9														1	1	1	Beam track
3	11														1	1	1	Entry fiducials
4	44														1	1	1	
5	XX															1		Short stopping protons
6	WW																	
7	YY																	Short stopping pions
8	ZZ																	
9	PP																1	Long stopping protons
10	AP																1	
11	OO																	
12	AO																	
13	QQ														1			Long stopping pions
14	AQ														1			
15	RR																	
16	AR																	Short kinks
17	UU																	
18	U1																	
19	VV																	
20	V1																	Long kinks
21	AT																	
22	TT																	
23	T1																	
24	AS																	
25	SS																	
26	S1																	
27	A1														1	1	1	Ordinary secondary tracks
28	A2														1		1	
29	A3														1			
30	A4																	
31	A5																	
32	A6																	
33	A7																	
34	A8																	
35	MM																	V zeros
36	M1																	
37	M2																	
38	NN																	
39	N1																	
40	N2																	
41	22														1	1	1	Exit fiducials
42	33														1	1	1	Apex check
43	AA														1	1	1	

Note: M/C = Machine



Correspondence between bits 10 to 15 in each word of the array DTOP,
and machines 1 to 6

Figure 4.8

4.4.3 The measurement of an event

After the topology has been decoded and the various arrays in common are set by TOPOL for the start of a new event, the appropriate word in VIEW is set equal to 1, and POSTN is also set equal to 1. The programme then calls for the first measurement in view 1 which is the apex of the event.

When the apex measurement is accepted its coordinates are prepared for output onto disk (see Section 4.7), and the address of its coordinates within the coordinate bank is stored in PNTS. The programme then scans through the array DTOP (see Table 4.3) checking only the bit for the machine concerned, until the next "1 bit" is found. As this will be in the second word, POSTN will take the value 2, and the programme will call for the second measurement which is the beam track.

The coordinates of the beam track are stored in the array TRACK as they are received, until the END TRACK setting of the latching switches on the control panel is detected. The routine FIT (see Section 4.6) is then called to determine whether or not the track was well measured. If the track is accepted its coordinates are prepared for output (see Section 4.7), and the number of coordinates, and the address of the first pair of its coordinates within the coordinate bank, are stored in the array TRCKS.

The next bit found in DTOP for this machine will be in word 3. This corresponds to the first entry fiducial 11, and when its coordinates are received they are stored in the array FDCL. Likewise the coordinates of fiducial 44, which is the next measurement to be requested, are stored

in FDCL.

The programme then scans the array DTOP again, to determine which measurement is required next. If the next measurement is a long stopping proton, as in the machine 1 example shown in Table 4.3, then the corresponding value for POSTN would be 9. The programme would then request the measurement of PP, the stopping point of the proton, to be measured next. Its coordinate pair would be prepared for output and the address of its coordinates stored in PNTS.

POSTN would then take the value 10, and the programme would request the measurement of the proton stopping track, AP. During the measurement of this track its coordinates are handled in the same way as the beam track coordinates were. That is, they are stored in the array TRACK as they are received, and when the track measurement has been completed the routine FIT is called to check on the accuracy of the measurements. Provided the measurements are acceptable the coordinates are prepared for output and the usual references are stored in TRCKS.

The programme continues to call for all the required measurements in this way, preparing the coordinates for output as they become available, and storing the necessary referencing information in PNTS and TRCKS. Finally the measurements of the exit fiducials are requested. Their coordinates are stored beside those of the entry fiducials in the array FDCL. It was hoped to have been able to make a calculation at this stage on the four fiducial measurements to check on their relative positions, but this has not been possible because of lack of core space.

Before the view is terminated the operator is requested to return to the apex so that a check can be made on the correct counting of the digitisers. The value read in is checked against the setting value which was given at the start of the view. If this is acceptable the coordinates of the four fiducials are prepared for output, the addresses of their coordinates within the coordinate bank are recorded in the array FDCLS, and the view is terminated.

The same sequence is repeated for views 2 and 3. After the measurement of view 3 is completed it only remains for all the referencing information in FDCLS, PNTS and TRCKS to be prepared for output (see Section 4.7).

4.4.4 Overlays

The possibility of using overlays to increase the effective core size of the computer was considered. The principle involved is to make two or more routines share the same core locations, and to read in from the disk whichever is required at any particular time. Whenever it is necessary to change the routine currently in core the programme has to wait while the new one is read from the disk.

During on-line measuring the disk access mechanism is usually positioned over the last sector of data to have been written, and will have to move across the disk if it has to fetch an overlay routine. As more and more data are written on the disk the access arm has to move further and further to reach these routines when they are required. For a disk only half filled with data the time involved would be about half a second. The next time the programme is ready to write more data onto disk there

will be a similar time delay as the access mechanism returns to the position required for writing more data. Since each operator can raise interrupts at the rate of more than one a second, delays of this order will become noticeable when six machines, or even four machines are being operated simultaneously.

The amount of time spent in moving the access mechanism can be kept to a minimum if the routines chosen for overlaying are carefully selected. They should, if possible, be routines which are only seldom used. They must also be sufficiently long so that the resulting saving of core space is more than equal to the space required for the extra instructions which are required to handle the overlays.

The only routines in the on-line system which are not used frequently are those whose main function is concerned with the termination of the programme. These routines are however, quite short, and could not of themselves provide the means of gaining extra space in core. Besides the terminating routines there are those which are used when an operator signs on or signs off, but again these are not long enough.

The routine TOPOL, which is used only at the start of each event could form the basis of an overlay system. Other routines would also have to be selected to share the same core locations and be brought into core only when required. Since practically all the other routines are frequently used throughout the measuring of each event, frequent movement of the access mechanism would still be required and the average time to service an interrupt would inevitably increase.

The choice then is between maintaining the rapid response time, or increasing the effective core size. For the on-line system a rapid response to each interrupt is fundamental to its successful operation, and so it was decided not to use overlays but to make the best possible use of the core space available in other ways.

4.5 Control of data transfers through the interface

The routine OEMI which services the level 3 interrupts, handles all movement of data through the interface. This involves the determination of the cause of each interrupt, and the reading in and writing out of data.

In the input mode the required data word is specified by a coded address before it is read by the computer. Similarly, in the output mode, the required lamp display is specified by a coded address before the data word is output by the computer.

Since the interface is connected to the computer through an OEM channel, all communication between the computer and the interface must use the different registers and control lines of the OEM channel. These are controlled by the programme in the CPU and so the operation of the interface is also controlled by the CPU.

4.5.1 Details of the OEM channel

(a) Input and output registers

All data coming into the computer must be transferred through the 16-bit input register. A read instruction loads the contents of this register

into a specified location in core storage.

The 16-bit output register is used to output data to the lamp displays of the measuring machines, and to transfer addresses to the interface to specify which lamp display is required, or which data word is required to be read in. Data are written into this register by means of a write instruction.

(b) Interrupt register

Whenever any of the six bits in the interrupt register (RI) is set by the interface, a programme interrupt on level 3 is caused by hardware control. The different bits are used depending on the cause of the interrupt (see Figure 4.9). In response to each interrupt the programme examines the interrupt register to determine the cause of the interrupt.

RI/1 is used during data transfers into the computer. It is set by the interface when data have been transferred to the input register and are ready to be read into the computer by a read instruction.

RI/2 is used during data transfers from the computer to the interface. Data from the computer are transferred to the output register by a write instruction. The interface informs the computer when it has accepted the data by setting RI/2.

RI/3 is not used.

RI/4 is used for interrupts from the core store of the cosmic ray experiment. It is set whenever the core store is ready to transfer its data into the computer.

RI/5 is used by the measuring machines. Whenever the ACCEPT DATA or the EXECUTE interrupt button is pressed on any of the machines, RI/5 is set.

RI/6 is used for 'forced interrupts'. It is set in response to the setting by the programme of control line 4.

The different interrupts are all set by hardware control. They are reset when a sense-and-reset instruction is given by the programme. The same instruction transfers the contents of the interrupt register into bits 0 to 5 in the accumulator as shown in Figure 4.9. A simple sense instruction may also be given, which reads the interrupt register without resetting it.

The interrupt register is inhibited when the AUTO/MANUAL switch on the monitoring panel within the interface is set to MANUAL (see Section 3.3.4).

(c) Control lines

The OEM channel has four control lines which may be set and reset by the programme. They are used to inform the interface of the function it is required to perform. These functions are shown in Table 4.4.

Control line 1 is used with a write instruction to inform the interface that an input address is available in the output register. The address may refer to the device status word within the interface or to one of the seven data words of one of the six measuring machines, or it may refer to an address within the core store (see Section 4.5.2). When the interface receives this address it transfers the contents of the address into the input register, and the sets RI/1 and RI/2 simultaneously to inform the computer that the required data may be read, and that the contents of the output register have been received.

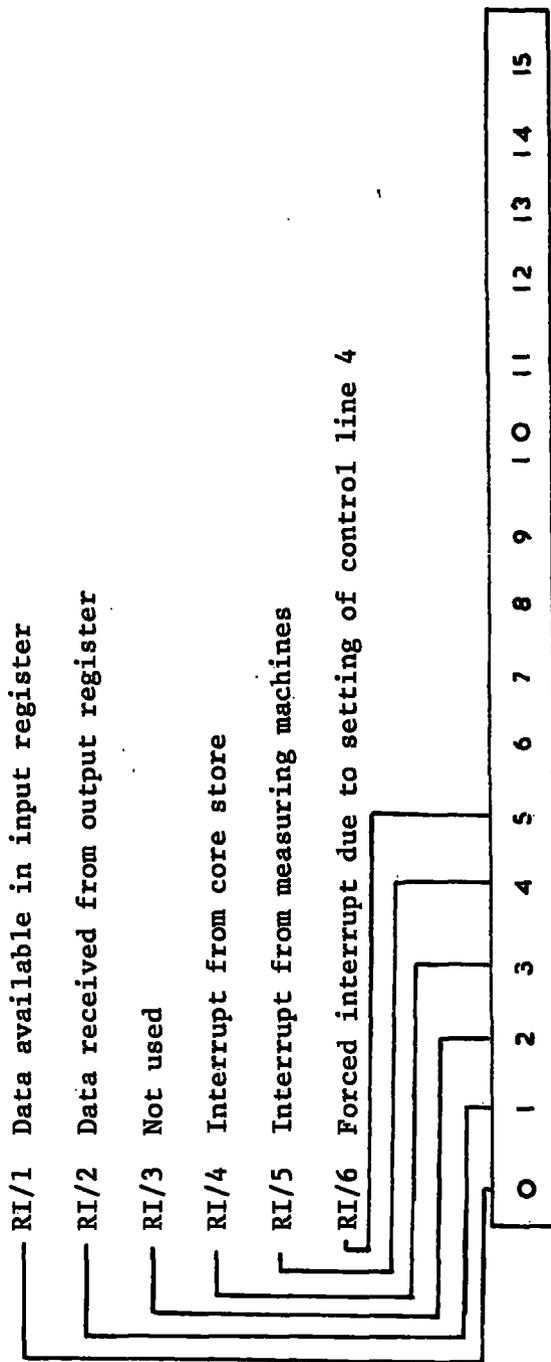


Figure 4.9 Interrupt register

Control line 2 is used with a write instruction to inform the interface that an output address is available in the output register (see Section 4.5.2). In this case the address will refer to the lamp display on one of the six measuring machines. When the interface has received this address it informs the computer by setting RI/2. Then, when the data word for the lamp display is output by the computer, the interface will transfer it to the machine indicated by the address that was given.

Table 4.4

Control Lines

CONTROL LINE	FUNCTION
1	Input address in output register
2	Output address in output register
3	Data for output in output register
4	Force interrupt by setting RI/6

Control line 3 is also used with a write instruction. It informs the interface that data are available in the output register. The interface transfers these data to the machine specified by the address that was given when control line 2 was set. The interface informs the computer when it has received the data by setting RI/2.

Control line 4 is used to raise a 'forced interrupt'. The interface responds by setting RI/6.

Each control line is reset by the programme after the interface has performed the required function.

4.5.2 The Westinghouse interface

(a) The Westinghouse device status word

By examining the interrupt register the programme is able to determine the cause of each interrupt. In the case of an RI/5 interrupt, the programme has still to determine which machine caused the interrupt, and whether the ACCEPT DATA or the EXECUTE interrupt button was used.

This information is held in a special register within the interface called the Westinghouse device status word (WHDSW). Twelve bits are used, one corresponding to each of the twelve interrupt buttons on the six measuring machines (see Figure 4.10). Whenever any of these is pressed, the corresponding bit in the WHDSW is set. For each machine the ACCEPT DATA bit is not reset until the fourth data word from that machine is read by the computer. The EXECUTE bit is reset when the Y coordinate is read in.

The WHDSW may be read by the programme in the same way as any data word, by specifying the first of the 44 input addresses used by the interface.

(b) The digital input and output words

Input of data to the computer from the measuring machines is in the form of 16-bit binary words. Figure 4.11 illustrates the seven digital input words that are associated with each machine.

The first four words correspond to the fifteen decade edge switches on the control panel. Each decade switch is in binary coded decimal and

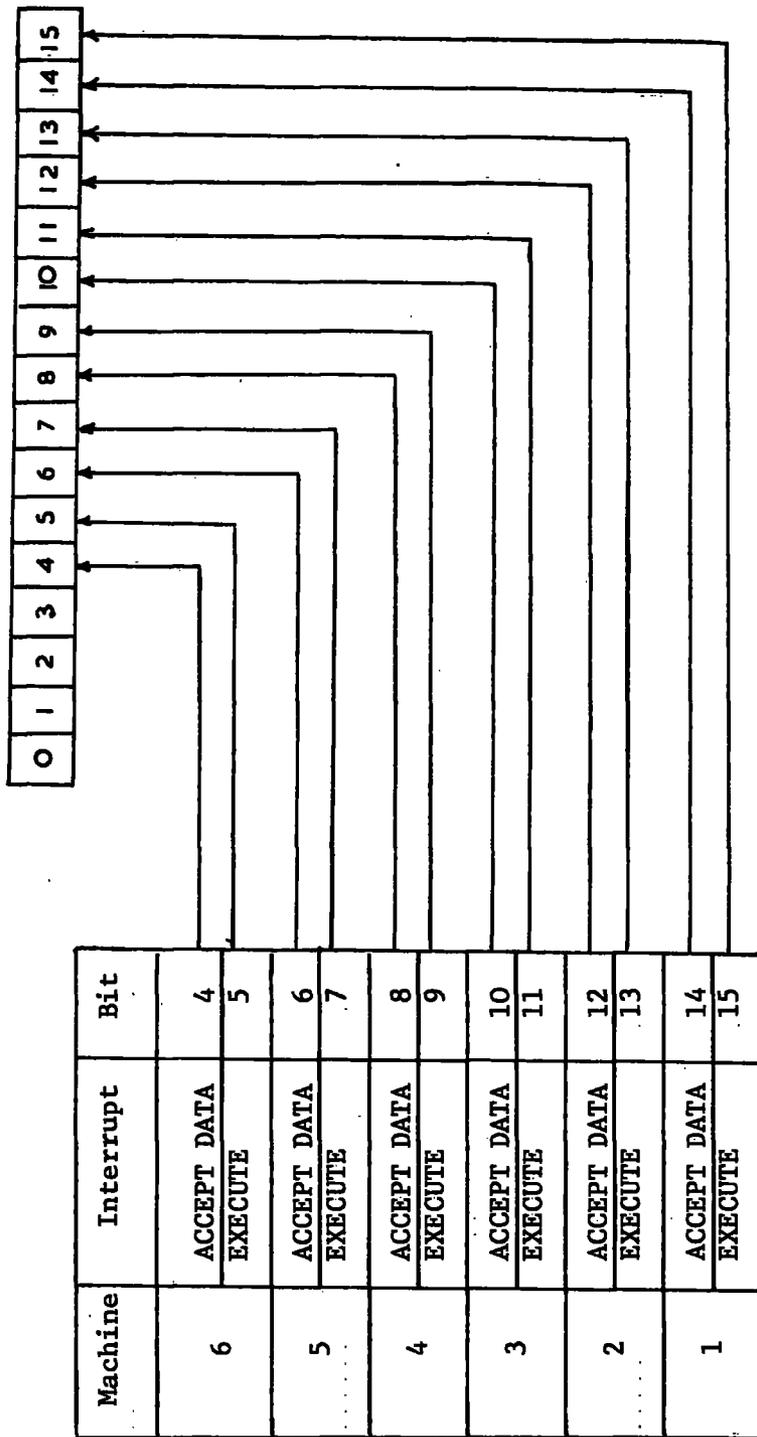
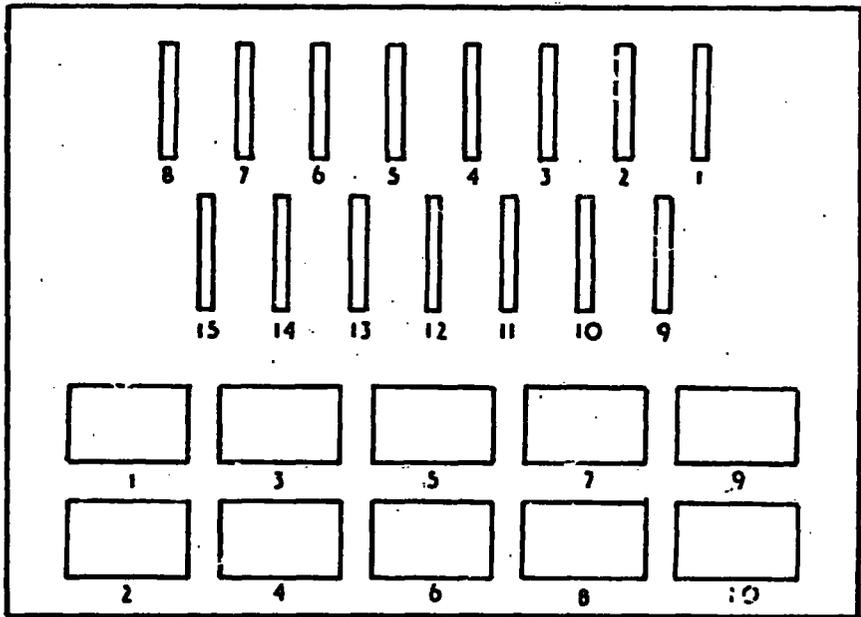


Figure 4.10 Westinghouse device status word



DECADE
EDGE
SWITCHES

ILLUMINATED
LATCHING
SWITCHES

CONTROL PANEL SECTION 2

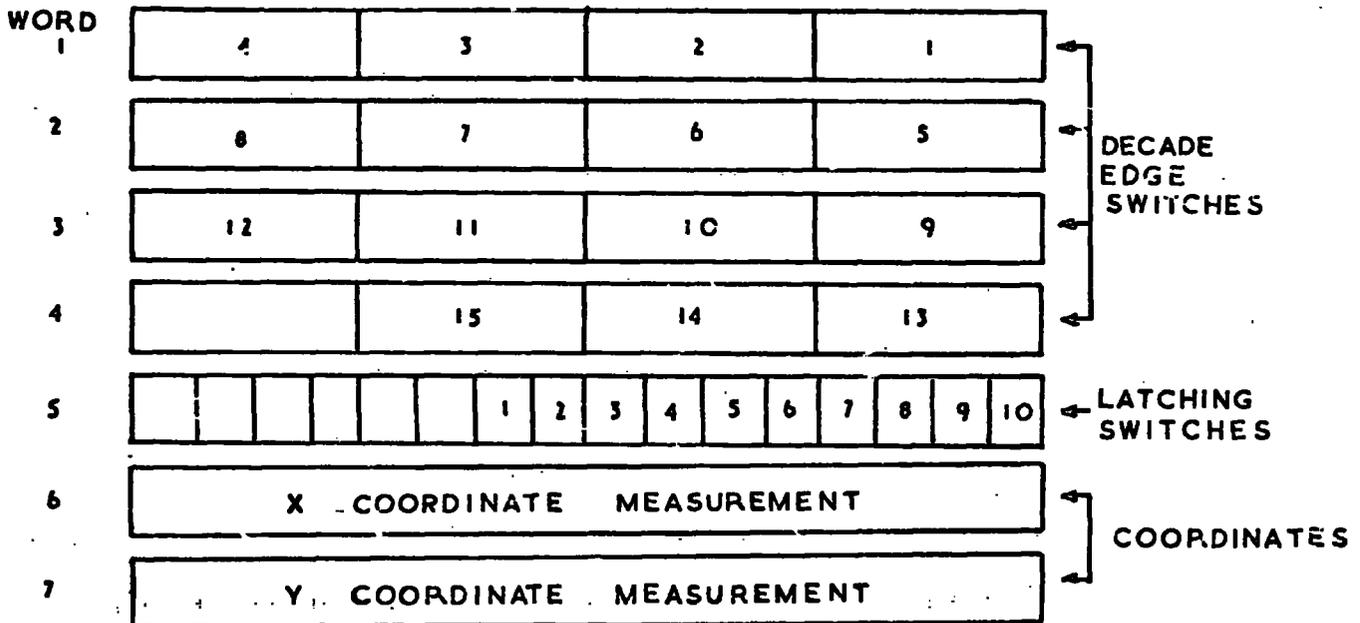


FIGURE 4.11 - DIGITAL INPUT WORDS -

so requires four bits. Only ten bits of the fifth word are used, one corresponding to each of the ten latching switches. The final two digital input words are the X and Y coordinate buffers respectively.

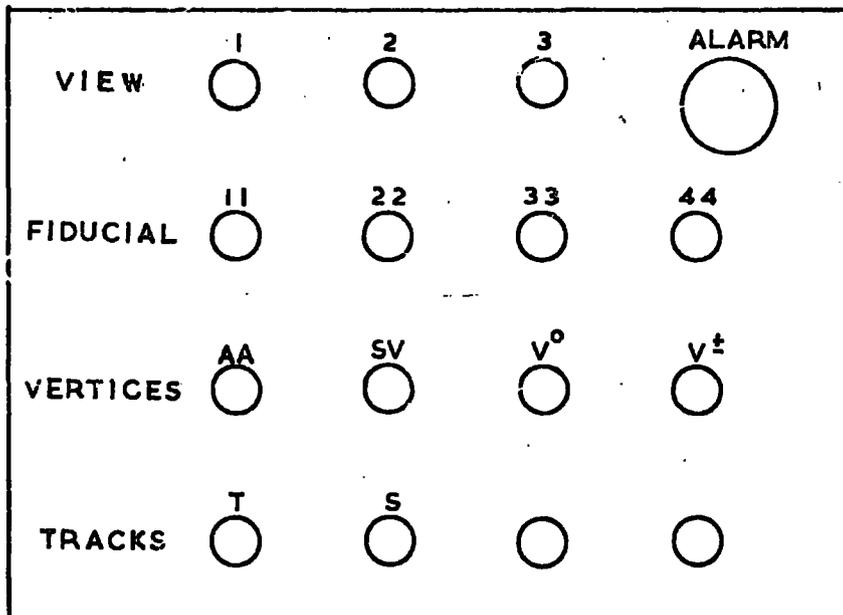
The output of data is also in the form of 16-bit binary words. A single digital output word is all that is required to illuminate the lamps on the control panel of a machine. This word may also be required to sound the alarm. Figure 4.12 shows the correspondence between the sixteen bits of the DOPW and the different lamps on the control panel.

(c) Input and output addresses

In order that the programme may be able to specify which data word it wishes to read in at any time, an address is assigned to each input word. These are listed in Table 4.5. The programme first of all gives the address corresponding to the word required. The interface decodes this address and transfers the appropriate data word into the input register of the OEM channel where it may be read by the programme.

An examination of Table 4.5 shows that the coded addresses corresponding to the seven input data words of each machine, are seven consecutive numbers. This makes it particularly easy for the programme to control the input of the seven data words from any of the six machines.

In the original specification for the interface, however, the seven addresses for one of the machines, machine 5, were not in the form of a sequence of consecutive numbers. This would not have been difficult to handle in the programme, but it would have required a longer and more complicated routine to control the input of data. The addresses given



CONTROL PANEL SECTION 1

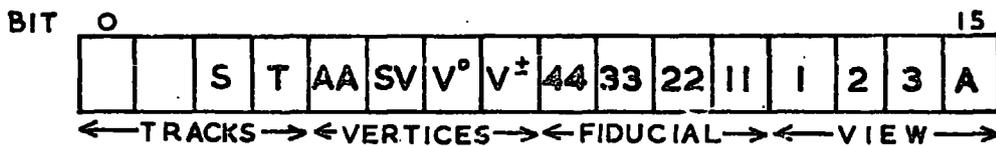


FIGURE 4.12 — DIGITAL OUTPUT WORD —

Table 4.5

Input Addresses

NUMBER	CODED ADDRESS	DEVICE REFERENCED
1	0010001	Westinghouse devise status word
2	0010010	Machine 1 Edge switches 1-4
3	0010011	" " " 5-8
4	0010100	" " " 9-12
5	0010101	" " " 13-15
6	0010110	" 10 latching switches
7	0010111	" X coordinate
8	0011000	" Y coordinate
9	0011001	Machine 2 Edge switches 1-4
10	0011010	" " " 5-8
11	0011011	" " " 9-12
12	0011100	" " " 13-15
13	0011101	" 10 latching switches
14	0011110	" X coordinate
15	0011111	" Y coordinate
16	0100001	Machine 3 Edge switches 1-4
17	0100010	" " " 5-8
18	0100011	" " " 9-12
19	0100100	" " " 13-15
20	0100101	" 10 latching switches
21	0100110	" X coordinate
22	0100111	" Y coordinate
23	0101000	Machine 4 Edge switches 1-4
24	0101001	" " " 5-8
25	0101010	" " " 9-12
26	0101011	" " " 13-15
27	0101100	" 10 latching switches
28	0101101	" X coordinate
29	0101110	" Y coordinate
30	1000001	Machine 5 Edge switches 1-4
31	1000010	" " " 5-9
32	1000011	" " " 9-12
33	1000100	" " " 13-15
34	1000101	" 10 latching switches
35	1000110	" X coordinate
36	1000111	" Y coordinate
37	1001000	Machine 6 Edge switches 1-4
38	1001001	" " " 5-9
39	1001010	" " " 9-12
40	1001011	" " " 13-15
41	1001100	" 10 latching switches
42	1001101	" X coordinate
43	1001110	" Y coordinate
44	1001111	Core store of cosmic ray experiment

in the original specification were therefore modified so that a set of seven consecutive numbers was used as input addresses for each of the six machines.

A similar set of addresses is used to refer to the lamp displays of the six measuring machines. These are illustrated in Table 4.6.

Table 4.6
Output Addresses

NUMBER	CODED ADDRESS	WORD ADDRESSED
1	0000001	Machine 1 lamp display
2	0000010	Machine 2 lamp display
3	0000100	Machine 3 lamp display
4	0001000	Machine 4 lamp display
5	0010000	Machine 5 lamp display
6	0100000	Machine 6 lamp display

During output of a DOPW the programme first of all specifies which lamp display is required by giving the appropriate address. The interface then opens the lines to that control panel so that when the DOPW is output by the computer it is routed by the interface to the required lamp display.

(d) The core store

Table 4.5 shows that the input address 44 is used to reference the core store of the cosmic ray experiment. Once this address has been given, all subsequent addresses refer to locations within the core store, until

address 1024 is given. This completes the transfer of data from the core store.

4.5.3 Programme instructions

(a) Input/output control commands

The OEM channel is operated under direct control from the CPU. Each command to the OEM channel is in the form of an XIO (execute input/output) instruction in conjunction with the appropriate IOCC (input/output control command).

Each IOCC consists of two consecutive words, the first of which must be in an even location. The first five bits of the second word contain the device code 11000, which specifies the OEM channel, and the following three bits specify the function that is required. The different functions that may be used are now described and summarised in Table 4.7.

<u>Write (001)</u>	word 1	core storage address
	word 2	1100000100000000

This command causes the contents of the core storage location specified to be transferred to the output register of the OEM channel.

<u>Read (010)</u>	word 1	core storage address
	word 2	1100001000000000

This command causes the contents of the input register of the OEM channel to be transferred to the core storage location specified.

<u>Control (100)</u>	(a) Set	word 1	control word
		word 2	1100010000000010

This command (with bit 14 equal to 1) is used to set the control lines.

This command, which has bit 15 equal to 1, causes the contents of the interrupt register to be transferred to bits 0 to 5 in the accumulator. At the same time the bits in the interrupt register are all reset. Once again bit 9 is used to specify that the instruction refers to the interrupt register and not to the other DSW which is not used in the on-line system.

If it is required to specify the other DSW then bit 9 in the sense instructions should be equal to zero.

(b) Control sequences for input and output of data

When it is required to read in data from the interface, the computer must first inform the interface which input word is required by giving it the appropriate address. This is done by setting control line 1 and transferring the address to the output register.

The interface decodes this address and transfers the contents of the address into the input register. About 20 μ sec later it sets RI/1 and RI/2 simultaneously. They inform the programme that the address in the output register has been received, and that the required data word is available in the input register. The programme then reads the contents of the input register, and resets control line 1.

The complete sequence for reading in a data word, and the corresponding Assembler language coding is shown in Figure 4.13.

A similar sequence of instructions is required to output a word to the lamp display of one of the measuring machines. In this case control lines 2 and 3 are both used.

Table 4.7

Input/Output Control Commands
for OEM Channel

Write (001)

core storage address	11000	001	00000000
	↓	↓	
	device code		function

Read (010)

core storage address	11000	010	00000000
----------------------	-------	-----	----------

Control (100) (a) Set

control word	11000	100	00000010
--------------	-------	-----	----------

(b) Reset

control word	11000	100	00000001
--------------	-------	-----	----------

Sense (111) (a) Sense only

	11000	111	01000000
--	-------	-----	----------

(b) Sense and reset

	11000	111	01000001
--	-------	-----	----------

	XIO	SET1	Set control line 1
	XIO	WRITA	Write input address in output register
WAIT	XIO	SENSE	Wait
	BNN	WAIT	for RI/1
	SLA	1	and RI/2
	BNN	WAIT	to be
	XIO	RESET	raised
	XIO	READ	Read data word from input register
	XIO	RSET1	Reset control line 1
	BSS	E 0	Ensure even location for IOCC's
SET1	DC	/8000	IOCC to
	DC	/C402	set control line 1
WRITA	DC	INAD	IOCC to
	DC	/C100	write input address
SENSE	DC	0	IOCC to
	DC	/C740	sense interrupt register
RESET	DC	0	IOCC to
	DC	/C741	sense and reset interrupt register
READ	DC	DIPW	IOCC to
	DC	/C200	read digital input word
RSET1	DC	/8000	IOCC to
	DC	/C401	reset control line 1
INAD	DC	x	Coded input address
DIPW	BSS	1	Location into which DIPW is read

Figure 4.13 Assembler Coding Required to Read in one Data Word

Note: x is coded input address corresponding to the data word required.

e.g. x = /001E if X coordinate from machine 2 is required.

The address of the required lamp display must be output first. It is written into the output register after control line 2 has been set. When the interface has received this address it raises interrupt RI/2. Again the response time is about 20 μ sec.

The programme then sets control line 3 and transfers the output word to the output register. When the interface receives this word it transfers it to the required lamp display, and raises RI/2 again.

The complete sequence for transferring a word to the lamp display on one of the control panels, and the corresponding Assembler language coding is shown in Figure 4.14.

(c) Timing considerations

The control sequence which is required to read in data from the interface includes a period during which the programme is waiting for the interface to raise RI/1 and RI/2. Similarly, during the output of data there are two occasions within the control sequence when the programme waits for RI/2 to be raised. In each case the time involved is about 20 μ s.

This time is so short that not more than about two instructions could be executed before the interrupt is raised. The most efficient way of handling interrupts within such a short interval is to execute the programme on level 3, which prevents the setting of the interrupt register from actually raising an interrupt. The programme can then sense the interrupt register repeatedly until the expected settings are observed.

When this method was used initially it was found that occasionally the programme failed to detect the setting of the interrupt register. At the

	XIO	SET2	Set control line 2
	XIO	WRITA	Write output address in output register
WAIT1	XIO	SENSE	Wait
	SLA	1	for RI/2
	BNN	WAIT1	to be
	XIO	RESET	raised
	XIO	RSET2	Reset control line 2
	XIO	SET3	Set control line 3
	XIO	WRITD	Write output word in output register
WAIT2	XIO	SENSE	Wait
	SLA	1	for RI/2
	BNN	WAIT2	to be
	XIO	RESET	raised
	XIO	RSET3	Reset control line 3
	BSS	E 0	Ensure even location for IOCC's
SET2	DC	/4000	IOCC to
	DC	/C402	set control line 2
WRITA	DC	OUTAD	IOCC to
	DC	/C100	write output address
SENSE	DC	0	IOCC to
	DC	/C740	sense interrupt register
RESET	DC	0	IOCC to sense
	DC	/C741	and reset interrupt register
RSET2	DC	/4000	IOCC to
	DC	/C401	reset control line 2
SET3	DC	/2000	IOCC to
	DC	/C402	set control line 3
WRITD	DC	DOPW	IOCC to
	DC	/C100	write DOPW
RSET3	DC	/2000	IOCC to
	DC	/C401	reset control line 3
OUTAD	DC	y	Coded output address
DOPW	DC	d	Digital output word

Figure 4.14 Assembler Coding Required to output one DOPW

y is the coded output address corresponding
to the lamp display required

d is the digital output word to be transferred
to the lamp display

time, the sense-and-reset instruction was being used to read the interrupt register, and it was deduced that the sense and reset functions could not have occurred exactly simultaneously. This meant that if RI/1 (or RI/2) was raised after the interrupt register was sensed and before it was reset, then it would not be detected by the programme.

The solution to this problem is to use the sense only instruction (C740) which does not reset the interrupt register. When the expected setting is observed the register may be reset by the sense-and-reset instruction (C741).

Once this change was made there was no further trouble from interrupts passing undetected.

4.6 Track checking

The curve fitting subroutine is used to check the coordinates of each track after the operator signals its completion with the END TRACK setting of the switches on her control panel. The purpose of this check is to ascertain whether or not the track was well measured.

It is not, however, possible to ensure that measurements which are accepted at this stage will also be accepted by THRESH. The reason for this is basically because only a two dimensional fit can be made during on-line measuring, and this cannot guarantee a successful three dimensional reconstruction. For instance, a track which is well measured in each of the three views will not reconstruct in three dimensions unless the measurements are indeed of the corresponding track in each view. Such a

check cannot be made at this stage. However, if two similar tracks belonging to one event are confused during the measurement of a view, then the measurements from the different views can be recombined correctly by the track matching facility in THRESH.

Since the tracks are formed by charged particles moving in a magnetic field, the path traced out by each particle will be that of a helix, provided there is no significant energy loss. As the cameras are mounted with their axes parallel to the magnetic field the tracks are viewed along the field direction and the helices will appear in each view as arcs of circles. For slower particles, where the loss of energy cannot be ignored, the paths will be more complicated. In projection however, these will still approximate to arcs of circles over short lengths of tracks.

Several different curve fitting procedures were considered, and some were tested using data from the paper tape system. These attempted to determine whether or not the coordinate measurements for each track lay along a smooth curve. In such a fitting procedure it is also advantageous to check that the fitted curve passes through, or at least very close to, the apex of the event. Unless this is the case the track will be rejected by THRESH.

It was also decided that a check on the sagitta of most of the tracks was desirable. This is because a sagitta of at least 2 mm in bubble chamber space is required for an accurate determination of the momentum of the track to be made. A track whose measured length has a sagitta of less than 2 mm is therefore rejected during on-line measuring. In the

case of a stopping track the momentum is calculated from the range energy relation and so no minimum value of the sagitta is required. Since all the beam tracks in a particular experiment have the same known momentum there is again no need to insist upon a minimum sagitta. The check on the sagitta then applies only to the ordinary secondary tracks which leave the chamber.

4.6.1 Circle fitting procedure

For the fast tracks which may be assumed to be circular on the film, a least squares fit to the measurements was considered, but this was found to be very difficult to do. However, it is relatively easy to fit a circle to a set of measurements by finding the values of the circle parameters g , f and c which minimise

$$\sum_{i=1}^n (x_i^2 + y_i^2 + gx_i + fy_i + c)^2$$

where (x_i, y_i) are the coordinates of the n points measured along the track, and the expression in brackets is the general equation of a circle.

This is justified by considering that, for any point (x, y) lying on a circle with parameters g , f and c

$$x^2 + y^2 + gx + fy + c = 0,$$

while for any point (X, Y) which does not lie on the same circle,

$$X^2 + Y^2 + gX + fY + c = v$$

where v is not equal to zero. It can be shown (see Appendix B) that, for points lying close to the circle

$$v = td \text{ (approx.)}$$

where t is the distance of the point (X, Y) from the circumference of the circle, and d is the diameter of the circle.

A true least squares fit to a circle would require the parameters of the fitted circle to be determined by minimising Σt^2 . Since small changes to the fitted circle would affect t rather than d , d may be considered to be constant, so that v is proportional to t to a first approximation, and the parameters of the fitted circle may be calculated by minimising Σv^2 instead of Σt , that is by minimising

$$\sum_{i=1}^n (x_i^2 + y_i^2 + gx_i + fy_i + c)^2$$

Having determined the parameters of the fitted circle it is then necessary to determine whether or not the fit is sufficiently good. Since for points lying close to the circle

$$v = td \text{ (approx.)},$$

the standard deviation may be calculated in the usual way by substituting v/d for t in the usual expression for the standard deviation. Thus

$$\text{standard deviation} = \frac{\Sigma t^2}{n_f} = \frac{\Sigma (v/d)^2}{(n - 3)}$$

where n_f is the number of degrees of freedom, and n is the number of points measured along the track, so that after a three parameter fit has been made

$$n_f = n - 3$$

The value of the standard deviation obtained in this way is used to determine whether or not the measurements of a track are considered to be

acceptable. At first the maximum acceptable standard deviation was fixed for all tracks. However, it was soon found that the lower momentum tracks were usually failing this test unless the measured length of the track was reduced by 50% or more. This is because these tracks were produced by particles which lost a significant proportion of their energy within the bubble chamber, resulting in tracks whose curvature was not constant, but increased noticeably along the track.

In THRESH this is handled by using a variable tolerance, which depends upon the curvature of the fitted circle. The same technique was adopted in the on-line system so that for tracks whose fitted circle has a diameter of less than 10^5 fringes, the maximum acceptable standard deviation is given by

$$\text{TOL} = A + B (1 - d/10^5)^2$$

where d is the diameter of the fitted circle in fringes.

For the high energy particles whose fitted circles have diameters of more than 10^5 fringes the second term in the above expression is not used, as these tracks should give a close fit to a circle. The highly energetic particles lose relatively little energy in the chamber, and so the curvature of their tracks is constant.

Initially the values of A and B were set so that the test on the track measurements was quite relaxed, but as the operators became familiar with the on-line system these values were gradually reduced to 8 and 40 respectively.

About 200 events have been measured with A and B set to more

stringent values of 4 and 20, but this had little effect on the pass rate (percentage of events which are successfully reconstructed in THRESH), and reduced the measuring rate by a factor of about three. It would appear that the normal values of 8 and 40 are sufficiently strict and that a better pass rate can only be achieved by including a test on the fiducial measurements.

4.6.2 Apex test

After the circle fitting procedure has been applied to the coordinate measurements of a track, the programme then determines whether or not the apex lies close to the fitted circle. This test requires very little calculation.

In the notation of Section 4.6.1 the distance of any point from the circle is given by

$$t = v/d$$

provided the point does not lie far from the fitted circle. Since the vertex is expected to lie close to the fitted circle the above approximation may be applied. It only remains therefore, to evaluate v/d for the apex.

Before the circle fitting calculation is made, the origin of the coordinate measurements is moved to the apex of the event, and so the evaluation of v/d for the apex requires the insertion of the coordinates (0,0) in the expression for v . This gives, in the notation of Section 4.6.1,

$$v_{\text{apex}} = c$$

so that

$$t_{\text{apex}} = c/d$$

The maximum acceptable value for t_{apex} is the same as that allowed for the standard deviation, and so depends upon the curvature of the fitted circle.

If a track measurement fails to satisfy either, or both, of these requirements, then the standard deviation and the distance of the apex from the fitted circle are printed out for the operator to see, and she is requested to remeasure the track.

4.6.3 Sagitta calculation

For beam tracks and stopping tracks only the circle fitting procedure and the apex test have to be satisfied. The measurements of all other tracks have also to satisfy a check on the sagitta.

As with the circle fitting and apex tests, only a two dimensional check on the sagitta can be made. This can only be a rough check on what the sagitta of the track would be when it is reconstructed in three dimensions. However, since the curvature of the track is due to the magnetic field, and the axes of the cameras are parallel to this field, the simple calculation made in two dimensions will give a close approximation to the true value of the sagitta in three dimensions.

The calculation of the sagitta is based upon the assumption that the first and last points measured on the track mark the extremities of the measured length of the track. The programme is written so that, if a track fails on account of its sagitta being too small, more points can be added without the whole track having to be remeasured. Clearly this will only work when the track is being measured outwardly from the vertex of the event.

For the track measurement to be acceptable the sagitta must be at least 70 fringes. This corresponds to about 2 mm in chamber space.

If the evaluation of the sagitta shows that an insufficient length of track has been measured, then a printed message is given to inform the operator of the extent of the failure. She is given the option of adding more points to those previously given, or of remeasuring the whole track.

If the measurements are accepted they are passed on to the output routines which prepare the coordinates for transfer onto disk, and preserve the necessary referencing information in the array TRCKS. The operator is then requested to proceed with the next measurement.

4.6.4 Precision

Since the track checking is an important part of the on-line system it was necessary to investigate whether or not the calculations should be done in extended precision. Standard precision arithmetic uses 23 bits for the mantissa, compared with 31 bits for extended precision. This corresponds to at most 7 and 10 significant decimal digits respectively.

When tests were made on data from the paper tape system, using both standard and extended precision, it was found that only minor differences in the third or fourth significant figure appeared in the evaluation of the standard deviation. Since it is only required to know whether or not this is less than the maximum acceptable value, the accuracy of the standard precision calculation is considered to be sufficient.

This means that less core space is required for the circle fitting and other arithmetic routines than would have been required had it been necessary to use extended precision, as all real variables in extended

precision require three words in core compared with two for standard precision.

The choice of standard or extended precision for real variables applies to all real variables in the programme. It is not possible to have some variables in standard precision while others are in extended precision. Having selected standard precision as being sufficient for the calculation of the standard deviation, it was then necessary to check that the same precision was sufficient for the calculation of the sagitta.

This was investigated by measuring the same length of a track several times over, and the value of the calculated sagitta printed out for each measurement. The same procedure was repeated on several tracks. In each case the values obtained for each track showed only slight variation. If serious rounding errors were affecting the calculation, then it would be expected that a random set of values would be printed out by this test. Since this was not found to be the case, it was accepted that standard precision was sufficient for the calculation of the sagitta.

4.6.5 Coordinate modifications

Before the circle fitting procedure is applied to the coordinate measurements certain modifications have to be made.

The coordinates are received from the measuring machines in the form of 16-bit words representing positive numbers in the range 0 to 65535. However, the 16-bit word when used as an integer in the 1130 is usually interpreted as a number in the range - 32768 to + 32767. This is achieved

by using bit 0, the most significant bit, as a sign bit, so that it takes the value - 32768 instead of + 32768.

Since the integral coordinates have to be converted to real (floating point) at some stage before the circle fitting procedure is applied to them, it was decided to take account of the positive value of bit 0 during this conversion. A special routine was written to convert the integral values to real, treating all 16 bits as positive. This required fewer words in core than the equivalent FORTRAN routine which would have added 655360 to each coordinate whose value was interpreted as being negative.

A further modification is applied to correct for the factor of two difference in precision between the X and Y measurements. The reason for this difference is explained in Section 3.3.2. If no correction were made for this in the programme, the tracks would have to be treated as arcs, not of circles, but of ellipses.

In Section 2.1.1 it was mentioned that on machines 1 and 2 the measurements in the Y direction are too small by about 8%. This is also corrected before the circle fit is made.

There are severe distortions on machine 4 which are described in Section 2.1.3. These are due to the angle of projection, and to the necessity of making measurements some distance from the optic axis of the projection system. Before the data are input to THRESH corrections are made for these distortions. However, since considerable core space is required by the routine which makes this modification, allowances for these distortions cannot be made during on-line measuring. The coordinates are corrected only

for the reduced precision on the X measurements. It is accepted that this is not satisfactory, but until the rest of the programme can be pruned sufficiently to provide enough space, there remains no option but to omit this correction. However, in spite of being unable to correct for these distortions before applying the circle fitting routine to these coordinates, the results obtained from machine 4 compare favourably with those obtained from the other machines.

4.6.6 Other curve fitting procedures

Section

The choice of the circle fitting procedure which is described in/ 4.6.1 was made after considering several other curve fitting procedures. These were examined and tested using data from the paper tape system.

A least squares fit to a parabola was tried. It was found that if a parabola of the form

$$y = ax^2 + bx + c$$

was used, then those tracks whose X coordinates passed through a maximum or a minimum value, could only be fitted after a rotation of axes. Then, after the coordinates have been transformed, the least squares parabola fit treated the coordinates as though the X values were perfectly accurate, while all the errors were in the Y measurements. The circle fit, on the other hand, requires no rotation of axes, and treats both the X and the Y coordinates in exactly the same way. Since there are errors on the X measurements as well as on the Y, and since the tracks are close approximations to arcs of circles, not parabolae, it was decided to use the circle, rather than the parabola fitting procedure.

Some investigation of an ellipse test was also made. This examined

the position of each point in relation to its immediate neighbours and had the advantage of being able to check each point as it was measured. This involved determining whether or not the second of the three points lay within an ellipse which had the other two points as foci. The choice of a suitable eccentricity for the ellipse was very difficult to make as the relative positions of the measured points depended on the curvature of the track and the spacing of the points along it. The test is in fact quite inadequate as it takes no account of whether the curvature of the track is positive or negative, and so would accept points which describe a track of continually varying curvature.

The possibility of using a more sophisticated circle routine than that described in Section 4.6.1 was considered. It was hoped that by examining the measured points on a track that had failed, it might have been possible to obtain an acceptable standard deviation by rejecting one or two of the measured points. However, it was found that if one point was very badly measured, then it was not necessarily true to say that the offending point was the one which lay furthest from the fitted circle. The only reliable way of determining which point or points to reject was found to be by removing each of the n points in turn, and applying the circle fitting procedure to the remaining $(n-1)$ points. This could require 3 or 4 seconds of CPU time during which all the machines would have to wait, and so it is not suitable for inclusion in the on-line system. In any case most of the failures cannot be corrected simply by removing one or two of the measured points. Failures of this kind are usually produced by failures on the measuring machines, whereas the majority of track failures are due to

inaccurate measuring on the part of the operators.

4.7 Output of data onto disk

The format used for the output of data from the measuring machines is based upon that required for input to THRESH.

Actual transfer of the data onto disk is carried out by the IBM supplied DISK1 routine. This was chosen as being the shorter of the two disk subroutines intended for use by programmes written in Assembler language.

4.7.1 Organisation of disk storage

Each disk has space on it for 500K words which are grouped into cylinders and sectors as shown in Table 4.8.

Table 4.8

Organisation of disk storage

<u>No.of</u> / <u>per</u>	<u>Sector</u>	<u>Cylinder</u>	<u>Disk</u>
Words	320	2,560	512,000
Sectors		8	1,600
Cylinders			200

The eight sectors which comprise one cylinder may be read or written without moving the disk access mechanism.

Usually about one third of the space on the disk, consisting of the IBM system area and the user area, is taken up by the disk monitor system,

compilers and various programmes including those for the on-line system. The rest of the space, about 1000 sectors, is called the working storage and may be used for data.

The data are written on the disk starting at the beginning of working storage. This was done originally because the programmes which handle disk input to NUMAC require the data to be stored there. Since the working storage is available to the disk monitor system, compilers and user written programmes as temporary disk storage, great care has to be taken to ensure that any disks which have data in working storage are not made available for general use, as otherwise some of the data could inadvertently be overwritten.

In order to make full use of the space available, the data must be written in complete sectors of 320 words each. It is for this reason that the space in core which is allocated to the disk output buffer is 320 words long.

4.7.2 General considerations

In order to make the output from the on-line system compatible with the required form of input to THRESH, the data are arranged in the form of 80 column card images. The programme which transmits the data from the disk into the IBM 360/67 in Newcastle requires the card images to be stored in working storage in packed EBCDIC (Extended Binary Coded Decimal Interchange Code), and so the on-line programme was written to output the data in this form.

When it was decided to edit the data before they were input to NUMAC there was no longer any particular advantage in using this format, except that the routine to do this was already working, and the data have to be arranged in this way at some time before they can be input to NUMAC. Some saving of core space is achieved by preparing the data in console printer code rather than in EBCDIC, since the console printer code routines are already required to handle the preparation of messages for the console printer. The sorting routine later translates the data from console printer code to EBCDIC as required.

Like EBCDIC the console printer code can be packed two characters per word and so each 80 column card image requires only 40 words. The 320 word disk output buffer can therefore accommodate eight card images.

Since the on-line system will control up to six measuring machines simultaneously, there will usually be data to output from several machines at a time. It was thought at first that the space on the disk could be divided into distinct areas, one for each of the machines being used. The data from each machine would then be transferred to its own data area. However, this would have required a separate 320 word output buffer for each of the machines, and there is not sufficient core space available for this. It would also have produced a timing problem as each time one of the disk buffers was full the programme would have had to have wait until the disk access mechanism had moved to the correct data area before transferring the data onto the disk. During this time, which could have been half a second or more, all the measuring machines would have had to have waited.

An alternative method is to use one disk output buffer. As the data from each machine become available they are transferred to the disk buffer in the form of card images which are tagged by machine number for later identification. When eight card images have been received, the contents of the buffer are written onto the next available sector on the disk. By writing sequentially there is no noticeable delay due to the movement of the access mechanism of the disk unit.

The programme does however have to wait while each sector is being written. The only way of avoiding this is to use the method of double buffering which allows one buffer to be prepared while the other is being transferred to the disk. This requires not only two 320 word buffers, but also some extra programming to change from one buffer to the other. Because of the need to economise on core space it was decided not to use double buffering, but to allow the programme to wait during each disk transfer. Since the data are being transferred to sequential sectors on the disk, eight sectors can be written without moving the access mechanism. When no movement is required, the average time required to write one sector is 30 ms. When the next sector is to be written on the next cylinder, the programme must wait a further 37 ms while the access mechanism is moved.

4.7.3 Preparation of data for output

Although only one disk output buffer is required in core, there are six separate card image arrays, one for each of the possible six machines. Each 40 word array, corresponding to an 80 column card image packed two characters per word, is tagged in column 80 with the number of the machine

to which it refers,

There is also a 72 word array which is used as a card image. It is found convenient to prepare the heading card for each event, and the cards containing the referencing information, in unpacked coding. Completed cards are converted into packed console printer code and stored in the first 36 words of the appropriate card image array. The complete 40 word card image, including the machine number in column 80, is then transferred to the disk output buffer in the usual way.

As coordinate measurements from each machine become available, they are stored in the appropriate 40 word card image array. Completed cards are transferred to the next available space in the disk output buffer. Incomplete card images are held until more coordinates are available to fill them.

When the disk output buffer has received a full complement of 8 card images, the complete 320 word buffer is transferred to the next available sector on the disk.

The first card image to be output for each event contains '444444' in columns 1 to 6. This is to provide THRESH with the necessary indicator at the start of each event. It is also used by SORT which recognises this card as marking the beginning of a new event. If a '444444' card is detected before the completion of an event, then the event must have been unfinished and so is discarded.

The '444444' card is set up in the appropriate card image array which already has the machine number in column 80. This 40 word array, corresponding to a complete 80 column card image, is then transferred to the next available space in the disk output buffer.

The '444444' card image is followed by the heading card which contains all the usual heading information. The format of this card is similar to that which was used in the off-line system, except for a minor modification which is required to include the 7 digit topology coding, instead of the 5 digits used previously.

The coordinate measurements of each track are held in an array called TRACK until the circle fit is made and accepted. The coordinates are then arranged in card image form in the appropriate 40 word card image array. When a full complement of six coordinate pairs have been received, the completed card is transferred to the disk output buffer. Measurements of single points, such as the apex or stopping points, are moved to the next available space on the card currently being set up for that machine. Fiducial measurements are held in core until the end of the view when all four will have been measured. They are then transferred to the appropriate card in the same way as all the other coordinates.

At the completion of the event any remaining coordinates are transferred to the disk output buffer even if this means that the last card does not contain a full complement of six coordinate pairs. The end of the bank of coordinates is indicated by a single card having 'SSSSS' in columns 1 to 6. This card is prepared in the same way as the very first card of the event which contains '444444' in columns 1 to 6, and is used to inform SORT that there are no further coordinates for the event and that the summary information follows immediately.

All the referencing information for the event is then arranged on card images in the format required by THRESH. Each card is set up in turn,

in the 72 word card array, in unpacked coding. It is then converted to packed console printer code and stored in the first 36 words of the appropriate 40 word card image array, in the same way as the heading card is prepared. The complete card image, including the machine number in column 80, is then transferred to the disk output buffer.

Following the last card containing referencing information, one further card image is output. This has 'ZZZZZZ' in columns 1 to 6, and marks the end of the data for that event. It is recognised by SORT as marking the end of an event. If this card is not found, or if the 'SSSSSS' card is not found, then the event is incomplete and is discarded by SORT.

4.7.4 Continuation addresses

The data are transferred onto the disk starting from the beginning of working storage. When the on-line programme is terminated, usually at the end of each shift, any data still in the disk output buffer are transferred to the disk. This usually means that the last sector to be written is only partially filled with data.

The address of the next available card image is preserved until the programme is restarted, usually at the beginning of the next shift. Three words are used for this. The first gives the address of the sector within working storage, the second gives the address of the first available word within the sector, and the third word confirms how many 40 word card images remain unused on that sector.

This information is stored on the disk in sector zero of working storage, which means that the first card image on that sector is not available

for data. The first three words are initially set to 0, 40 and 7 respectively so that the first time the disk is used, the data are transferred to the disk starting at word 40 of sector zero of working storage.

At the start of the next shift the information in the first three words of sector zero of working storage enables the programme to transfer the new data onto the disk immediately following the end of the data from the previous shift. In this way one continuous block of data is written onto the disk.

As each sector is written the programme checks to see how many sectors are still available. A warning message is printed when only 15 sectors are left. This allows the operators to complete the events that they are measuring, provided no more than 15 sectors are required for the rest of the data. This is more than sufficient as each event requires about 30 card images on average, so that 15 sectors should accommodate four complete events. When the operators have completed their events they should terminate the programme in the usual way, and replace the filled disk with a fresh one.

4.8 Loading and termination of the on-line system

4.8.1 Initialising and loading programme

Each time the on-line programme is loaded into core prior to execution, certain initialising procedures have to be carried out. As the part of the programme that does this is only required during loading it is not included in the main programme that resides in core during measuring, but is run as a separate programme.

This programme, called LREAP, begins by calling another routine INIT

which clears the whole of common and then sets the machine numbers in the appropriate words in the array cards. It then reads the first sector of working storage to find the address of the first available space on the disk for the output of data. The sector address enables INIT to check that there are at least 15 sectors still available for data. If the sector specified has some data on it already the programme reads the sector into core and sets the counters so that the first data to become available will be used to fill up this sector.

LREAP then calls another routine, OEM, which clears any interrupts from the measuring machines that may have been raised while the programme was not running, and turns off all the lights on the lamp displays.

In order that the messages which are given on the console printer should provide a useful log of the measuring, LREAP requests the computer operator to type in the date and the time. The operator is then requested to set the console entry switches to indicate which machines will be operated, and whether or not the cosmic ray experiment will also be running. Although these switches are set at the beginning of each shift, or whenever the programme is loaded, they may be changed at any time during the running of the programme.

The programme LREAP finally executes a link to MAINE which is the main programme in core during on-line measuring. This causes MAINE to be loaded into core. LREAP is overwritten, but the arrays in common which were initialised by LREAP, are preserved from one link to the next. Control is then transferred to MAINE and the system is ready to attend to the first interrupt.

4.8.2 Termination of the on-line programme

The closing down routine is simpler than the loading routine, and since it is quite short it is included in the main core load. The request for termination of the on-line programme is given by resetting console entry switches 1 to 6 and pressing PROGRAM START.

When MAINE detects that these six console entry switches are all off, it calls the terminating routine END, which transfers to the disk any data still in core which have been prepared for output. Such data will have already been set up on card images and stored in the disk output buffer.

END then reads from the disk the first sector of working storage, updates the first three words to give the address of the next available card image on the disk and rewrites the modified sector back onto disk. Finally END prints the message "END OF RUN" to inform the operator that the programme has terminated correctly, and returns control to MAINE.

Provided switch 15 is also off, indicating that the programme is not required to service interrupts from the core store of the cosmic ray experiment, MAINE terminates the programme with an EXIT instruction.

CHAPTER . 5

CONCLUSIONS

A brief comparison is made between the paper tape measuring system and the on-line system, and various suggestions are given of ways in which the latter system could be improved.

5.1 Comparisons

The main aims of the on-line system are to provide a satisfactory and versatile measuring system which overcome as many of the problems of the paper tape system as possible.

Clearly the first advantage of the on-line system is that all errors due to mispunching of the paper tape are eliminated. Thus the data are more reliable, and there is no longer any loss of measuring time due to the maintenance of the tape punching equipment. The Westinghouse interface and the OEM channel have proved to be so reliable that no measuring time has been lost because of the maintenance of these items.

The second advantage of the on-line system is that, since the data are transferred directly into the computer, it is possible to make an immediate evaluation of the accuracy of the measurements. This encourages careful measuring on the part of the operators, and also enables failures of the measuring machines to be detected as they occur. Again this results in more accurate data, and hours or days of useless measuring on a faulty machine can be avoided.

The on-line system also relieves the operator of much of the routine work which used to be involved in measuring. In particular she no longer

needs to precede each measurement with the appropriate label. The method of generating labels internally, on the basis of a coded topology which is entered at the start of each event makes for more rapid measuring, but it also reduces the apparent flexibility of the system as the operators have to follow a strict order of measuring. However, the experience of the paper tape system showed that this flexibility was not required, and so its absence from the on-line system cannot be considered to be a disadvantage.

The advantages of all these changes to the measuring system are reflected in the various statistics which are available. In considering these figures it should be remembered that the same measuring machines are being used for the on-line system as were used for the paper tape system, and that it is only the way in which the machines are used that is different.

The overall rate of measuring has increased by about 20%. This figure is obtained from machines 1, 2 and 3 for the measurement of 4 prong events. (Machine 4 is not included as little work was done on that machine under the paper tape system). Slower motors were put onto machine 3 when the digitisers were changed to provide coordinate measurements compatible with the on-line system. If, however, it had been possible to retain the speed of the old motors, so that the old measuring rate could have been maintained, then the average measuring rate for the three machines would have increased by about 50%.

The pass rate (percentage of measured events which are successfully reconstructed by THRESH) has also shown a marked improvement. The average pass rate for events measured in the paper tape system was 52%, compared with 74% for the on-line system. Of the 52% originally accepted by THRESH

a number have later been returned for remeasurement as their tracks were rejected by the sagitta check which was introduced later. Thus the pass rate for the paper tape system was in fact rather less than the 52% quoted above. Comparing this with the on-line system an increase of about 50% in the pass rate is produced.

These pass rates were obtained for 3 prong events each of which involved the measurement of 4 tracks (3 secondaries and the beam track). The corresponding pass rates for individual tracks were 85% for the paper tape system and 93% for the on-line system. It is anticipated that a check on the accuracy of the fiducial measurements will further improve the pass rate.

5.2 Suggested modifications to the on-line system

Although the on-line system has been quite successful in achieving what it was designed to do, a few suggestions are given of various ways in which the system could be improved.

5.2.1 Output of data

About one quarter of the 8K words of core are allocated to common. This includes two disk output buffers and a number of arrays, some of which are used for keeping a record of the progress made throughout the measurement of each event. Most of the space however, is required for the recording of data and for the preparation of the data in card image form for output onto disk. This format for the output was based, as explained earlier, on that required for input to THRESH. Since the data are not input to THRESH directly but are edited and sorted first, it is not necessary to prepare the data in card image form at this stage.

The data could be output in the simplest form possible, in binary, and the editing programme SORT would arrange the data on card images in the required way.

A further suggestion would be to tag the measurements not only with the machine number to which they belong, but with the view number and the label of the point or track to which they refer, so that it would no longer be necessary to hold in core all the referencing information which is used to identify the different measurements within the bank of unlabelled coordinates. The output could in fact be rather similar to that of the paper tape system, and the programme SORT would then have to rearrange the data and set up the card images in much the same way as REAP did for the paper tape system.

This would mean that about half of common would no longer be required, and much simpler and shorter routines could also be used to prepare the data for output. The extra space in core would immediately open up the possibility of including more routines to check on the accuracy of the measurements.

5.2.2 Inhibition of interrupts

Associated with each interrupt button is an interrupt lamp which is situated on the control panel, and which goes out when the interrupt button has been pressed. It does not come on again until the data from that machine are read by the computer. So long as the interrupt lamp is off no further interrupts can be raised by the corresponding interrupt button. If the operator continues to press this button while the lamp

is off, nothing happens.

Usually the response from the computer is so rapid that only a slight flicker of the lamp is noticed, and the operators quickly become accustomed to raising interrupts as fast as they can measure. There are occasions however, when there is a delay in the response to an interrupt and the interrupt lamp does not come on again immediately. If the operator presses her interrupt button again before the interrupt lamp comes on, she may not be aware that no interrupt has been raised as no indication of this is given.

The computer cannot be programmed to warn the operator when this happens, as the computer itself is unaware of the situation. However, it should be possible, by making a hardware modification, to make the interrupt button sound the alarm on the control panel when it is pressed while the interrupt lamp is off. Since the sounding of the alarm would not, in this instance, be accompanied by a printed error message, the operator would know that it was caused by her pressing the interrupt button while it was inhibited.

5.2.3 Interruption of curve fitting

The response to most interrupts is so rapid that only a slight flicker of the interrupt lamp can be observed, and the new request on the lamp display appears to be given instantaneously. However, the response to certain interrupts requires more time, and in these cases a slight delay can be detected. This is particularly noticeable at the end of each track when the curve fitting routine is called, and at the end of

view 3 when the event is terminated. In neither of these cases is this at all inconvenient to the operator, as in the first case she will require a few moments to find the next point or track to be measured, and in the second she will have to prepare to start the next event.

However, the longer time required to service these interrupts can be a cause of interference to the operators of any other machines that are being used at the same time. This is because the programme is never forced to abandon the servicing of one interrupt to attend to another, and so, if it is busy attending to an END TRACK interrupt from one machine, all the other machines will have to wait until the curve fitting is completed. In this way the time required for curve fitting affects not only the machine whose measurements are being fitted, but also all the other machines that happen to be interrupting at that time. The same argument applies to the situation which arises when an event is terminated.

It would therefore be useful in the case of track fitting and event termination, to allow another machine to interrupt so that only the operator of the machine concerned would notice the delay. Interrupts from the other machines would then be attended to as they occur. Some care would have to be taken in implementing this suggestion as the programme must not be required to abandon checking the measurements from one machine to do the same for another machine.

5.2.4 Overlays

The main reason for rejecting the use of overlays to increase the effective core size of the computer, was that it would cause too much of

a delay in the response to an interrupt. However, there are ways of reducing this problem so that it could well be possible to use overlays without seriously affecting the response time.

Since the system has two disk drives, it is possible to arrange for the data to be output onto one disk and have the overlay routines on another disk. The access mechanism on the first disk drive would be positioned over the data, and on the second drive over the overlay routines, so that very little movement of the disk heads would be required either for fetching an overlay, or for transferring data to the disk. Consequently the speed of response to an interrupt would be little affected.

So long as the cosmic ray experiment is not being run on-line to the computer while the measuring machines are being operated, this arrangement provides a most effective way of avoiding the timing problem which is associated with the use of overlays. However the solution is not so simple when the system is required to handle the collection of data from the core store of the cosmic ray experiment as these data are output onto the second disk and the problem of the time required by the disk access mechanism to move from one part of the disk to another is only transferred from one disk drive to the other.

Another way of keeping the time delays to a minimum is to use one disk, as originally planned, but to make sure that it never has very much data on it so that the disk access mechanism never has very far to move. This can be achieved by transferring the data to another disk whenever the programme is terminated.

5.2.5 Fiducial check

At present no check is made on the accuracy of the fiducial measurements. This is because there is not sufficient space in core to accommodate the required routines. However, if by changing the format in which the data are output onto disk more core space is made available, or if the effective core size is increased by the use of overlays, then it should be possible to evaluate the accuracy of the fiducial measurements.

This would require examining their relative positions and comparing them with the expected positions for the view concerned. Information relating to their expected positions would have to be available for each of the three views and for different experiments as the machines are frequently measuring film from different experiments. It might even be necessary to hold special values for machine 4 as the distortions produced by its projection system affect the measured positions of the fiducials.

5.2.6 Curve fitting tolerances

It was mentioned in Section 4.6 that initially the curve fitting test was quite relaxed, and that as the operators became familiar with the on-line system the test was gradually made more stringent. Although during normal measuring there is no need to vary the stringency of this test it would be helpful for the training of new operators to be able to relax the test for one or more of the machines when required without affecting the other machines.

At present the maximum acceptable tolerance is determined by two

constants within the programme and so is the same for all the machines. Thus it is not possible to provide a strict test for an experienced operator, while at the same time to provide a more relaxed test for training a new operator.

If, however, the maximum acceptable tolerance was entered by the operator during the signing on procedure, then each operator could improve the accuracy of her work at her own rate. New operators could be trained with a relaxed test while experienced operators could continue to measure under strict test conditions. A check on the tolerances used by each operator could be made by printing out the values entered as each operator signs on. Alternatively these values could be entered during the initialising procedure so that they could not be changed by the operator during her shift.

5.2.7 Digital output words

Finally a very minor modification is suggested, which concerns the use of the fifteen lamps on the control panel which are used to indicate to the operator what she should measure. The table of digital output words (Table 3.2) shows that there are several instances where the same pattern of lights is used for two different measurements. The same table also shows that two of the lamps are never used.

It would assist the operators if the digital output words were chosen so that no two consecutive measurements were called for with the same pattern of lights. This is particularly important in the case of ordinary secondary tracks. Table 5.1 shows a suggested new set of

digital output words which makes use of the two spare lamps to give the operators more information than they receive at present.

These two lamps are labelled N and P. The N lamp is used as an indicator which is off for the first of two similar measurements and is on for the second. In the case of the ordinary secondary tracks the P lamp is also used and forms a 2-bit binary counter with the N lamp. The P lamp could normally be used to distinguish between proton and pion stopping tracks.

5.3 Final conclusions

The on-line system produces a 50% increase in the measuring rate compared with the paper tape system, and a similar increase in the proportion of events that are successfully reconstructed by THRESH. Taking these two increases together, the overall increase in the number of acceptable events which are measured per machine per hour is over 100%.

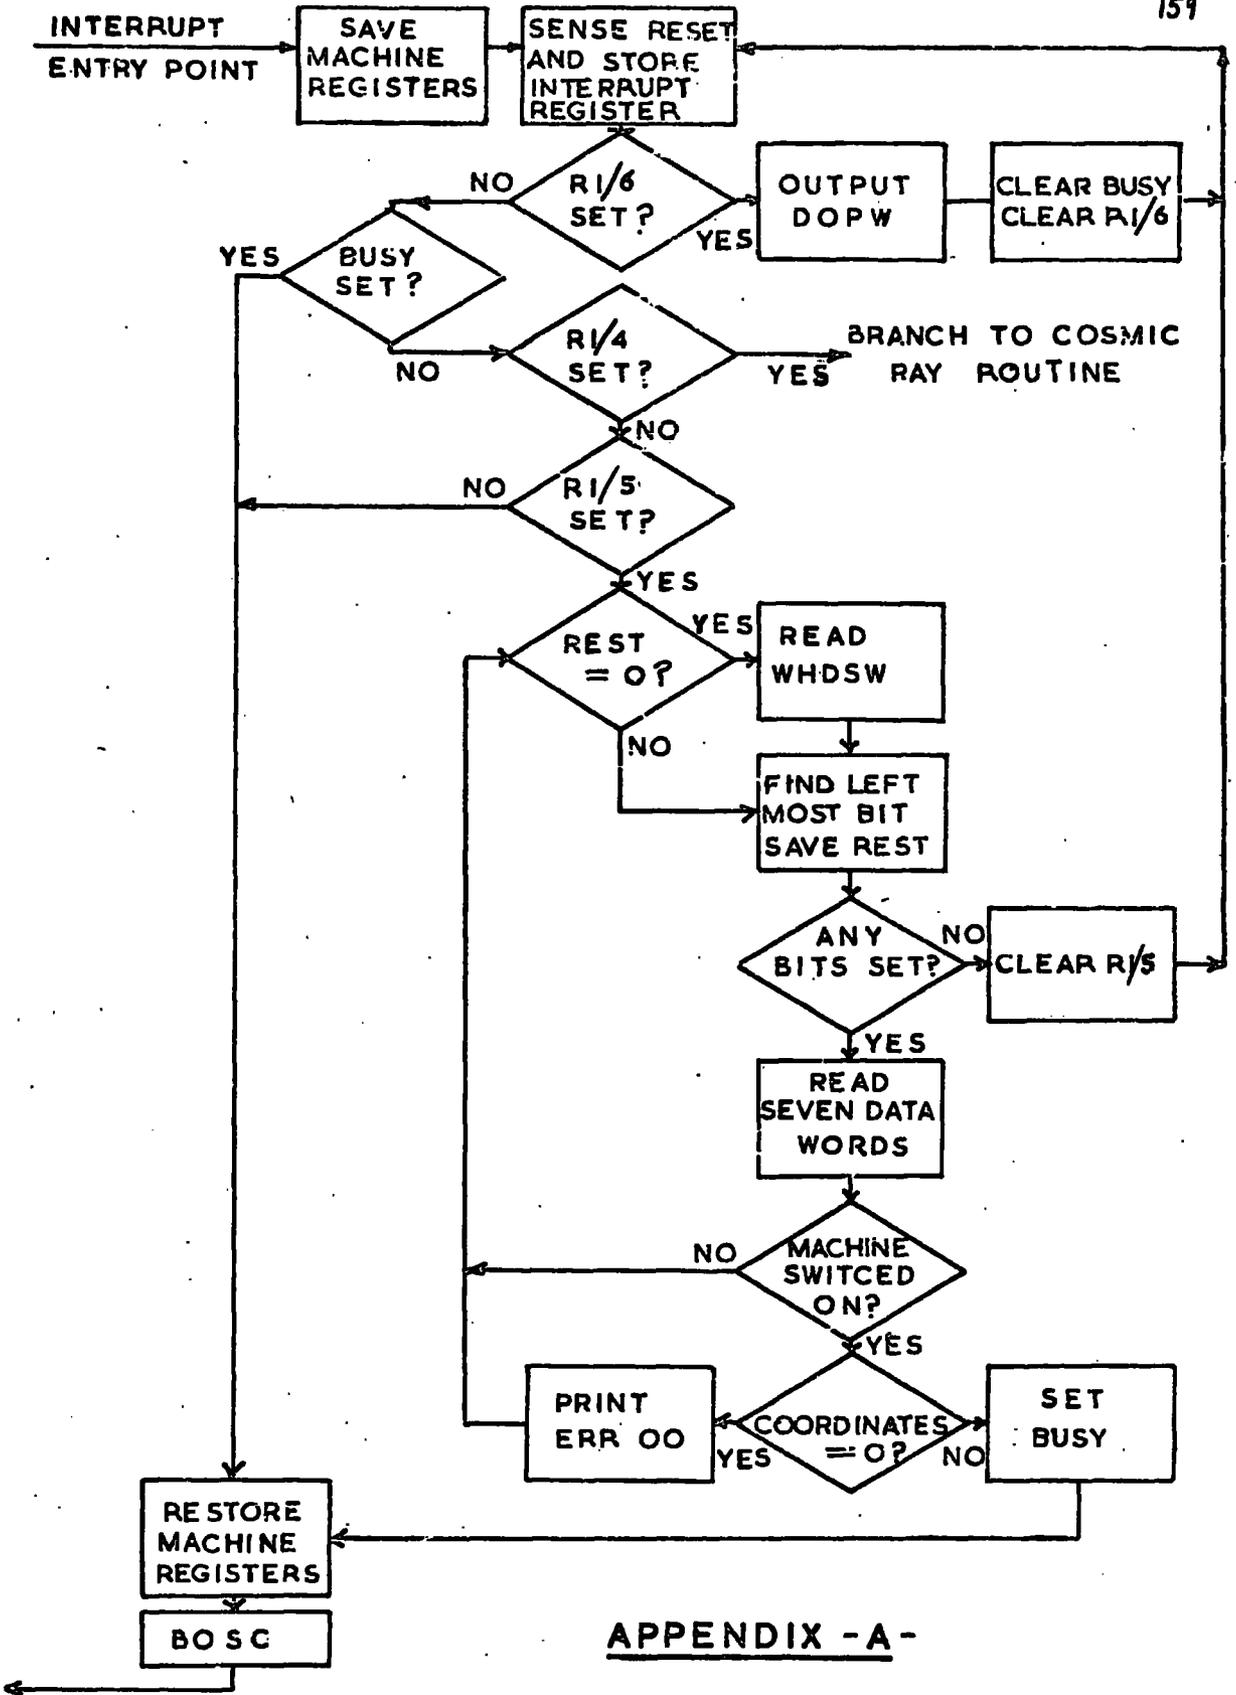
It is anticipated that this figure could be increased still further by implementing some of the suggestions which were given in Section 5.2.

Suggested new digital output words

NUMBER	LABEL	DIGITAL OUTPUT WORD											DESCRIPTION				
		TRACKS				VERTICES				FIDUCIAL				VIEW A			
		N	P	S	T	AA	SV	V ⁰	V ¹	44	33	22		11	1	2	3
1	AA					1											Apex
2	A9				1	1											Beam track
3	11												1				Entry fiducials
4	44									1							
5	XX		1				1										Short stopping protons
6	WW	1	1				1										
7	YY						1										Short stopping pions
8	ZZ	1					1										
9	PP		1			1	1										Long stopping protons
10	AP		1		1	1	1										
11	OO	1	1		1	1	1										
12	AO	1	1		1	1	1										
13	QQ					1	1										Long stopping pions
14	AQ				1	1	1										
15	RR	1				1	1										
16	AR	1			1	1	1										
17	UU								1								Short kinks
18	U1			1					1								
19	VV	1							1								
20	V1	1		1					1								
21	AT				1	1	1										Long kinks
22	TT					1	1										
23	T1			1			1										
24	AS	1			1	1	1										
25	SS	1				1	1										
26	S1	1		1			1										
27	A1				1	1											Ordinary secondary tracks
28	A2	1			1	1											
29	A3		1		1	1											
30	A4	1	1		1	1											
31	A5				1	1											
32	A6	1			1	1											
33	A7		1		1	1											
34	A8	1	1		1	1											
35	MM								1								V zeros
36	M1				1		1										
37	M2	1			1		1										
38	NN						1										
39	N1				1		1										
40	N2	1			1		1										
41	22											1					Exit fiducials
42	33											1					
43	AA					1											Apex check

Note: The appropriate view lamp is also illuminated

A P P E N D I C E S



APPENDIX -A-

- LOGIC DIAGRAM OF OEMI ROUTINE -

APPENDIX BThe circle fitting procedure

The equation of a circle may be written

$$X^2 + Y^2 + gX + fY + c = 0 \quad (1)$$

where g , f and c are the parameters of the circle. The centre of this circle is the point $(-g/2, -f/2)$, and its radius is equal to $\sqrt{g^2/4 + f^2/4 - c}$. Equation 1 holds good for any point (X, Y) which lies on the circle.

For any point (x, y) which does not lie on this circle

$$x^2 + y^2 + gx + fy + c = v \quad (2)$$

where v is not equal to zero. Equation 2 may be rearranged as

$$(x + g/2)^2 + (y + f/2)^2 = v + f^2/4 + g^2/4 - c \quad (3)$$

which in turn may be expressed as

$$p^2 = v + r^2 \quad (4)$$

where p is the distance of the point (x, y) from the centre of the circle, and r is the radius of the circle. From equation 4

$$v = (p + r)(p - r) \quad (5)$$

If the point (x, y) lies close to the circle so that $p = r$ (approx.), then equation 5 may be written

$$v = dt \text{ (approx.)} \quad (6)$$

where d is the diameter of the circle, and t is the distance of the point (x, y) from the circle.

If g , f and c are now considered to be the parameters of a variable circle which is fitted to the n points (x_i, y_i) , then modifying the above notation,

$$v_i = x_i^2 + y_i^2 + gx_i + fy_i + c \quad (7)$$

and

$$v_i = dt_i \text{ (approx.)} \quad (8)$$

where d is now the diameter of the fitted circle, and v_i and t_i refer to the point (x_i, y_i) . The approximation of equation 8 is valid since, for a fitted circle, all the points (x_i, y_i) must lie close to the circle.

The determination of the circle parameters g , f and c which correspond to the circle with the "best fit" to the points (x_i, y_i) , depends upon what is meant by the "best fit". A least squares circle fit would involve minimising $\sum t_i^2$, which from 7 and 8 above is equivalent to minimising

$$\sum \frac{(x_i^2 + y_i^2 + gx_i + fy_i + c)^2}{g^2 + f^2 - 4c}$$

When the partial derivatives of this expression with respect to g , f and c are set equal to zero to determine the turning value of $\sum t_i^2$, three rather complicated equations are obtained whose solution should give the required values of g , f and c . However these equations are not easily solved and an alternative method of determining the "best" values of g , f and c is preferred.

The approximation of equation 8, which was obtained from equation 5, effectively ignores differences in d which are of the order of t_i . That is, d may be considered constant relative to t_i , and so t_i may be considered to be proportional to v_i and a close approximation to minimising $\sum t_i^2$ may be obtained by minimising $\sum v_i^2$. Now

$$\sum v_i^2 = \sum (x_i^2 + y_i^2 + gx_i + fy_i + c)^2 \quad (9)$$

and when the partial derivatives of this expression with respect to g , f

and c are set equal to zero to determine the turning value of $\sum v_i^2$, three simultaneous linear equations in g , f and c are obtained.

$$\frac{\partial \sum v_i^2}{\partial g} = g \sum x_i^2 + f \sum x_i y_i + c \sum x_i + \sum x_i^3 + \sum x_i y_i^2 = 0 \quad (10)$$

$$\frac{\partial \sum v_i^2}{\partial f} = g \sum x_i y_i + f \sum y_i^2 + c \sum y_i + \sum x_i^2 y_i + \sum y_i^3 = 0 \quad (11)$$

$$\frac{\partial \sum v_i^2}{\partial c} = g \sum x_i + f \sum y_i + c \cdot n + \sum x_i^2 + \sum y_i^2 = 0 \quad (12)$$

where n is the number of points being fitted. The solution of these equations gives the values of the circle parameters g , f and c which correspond to the fitted circle. The fit obtained by this simple calculation is a very good approximation to a least squares circle fit.

APPENDIX CSORT

The data from the on-line system are rearranged and edited by a programme called SORT before being input to NUMAC for further processing.

(a) Input to SORT

During on-line measuring the data are output onto disk in the form of 80 column card images in console printer code packed two characters per word. Each card image is tagged in column 80 with the number of the machine to which it belongs. The data are output onto disk sequentially so that card images from each machine are interspersed with those of the other machines that were being operated at the same time. (see Figure C.1).

The data are output onto that part of the disk which is called working storage. Before they can be read by the programme SORT, the data have to be transferred to the user area where they can be referred to under the file name DATA1. Data stored under any other name are most easily made available to SORT by using a "*EQUAT" record.

(b) Output from SORT

After the data have been processed by SORT they are ready for input to NUMAC. They are stored on the disk in 80 column card images in packed EBCDIC. Only complete events are output by SORT and these are selected machine by machine. (see Figure C.2).

Before SORT is executed, the programme to handle the data as they enter NUMAC, together with all the necessary job control records, must be stored on the output disk starting from sector zero of working storage. These records are stored in card image form in packed EBCDIC. The first

Figure C.1 - Sample output from on-line measuring

30775 30336 28586 28752 32685 32749 32483 32697 31820 32521 30881 32258
29478 31831 27560 31159 25465 30319 23186 29266 20614 27910 32632 32729
36214 37492 24052 36706 44703 22612 36094 22488

SSSSSS

4
101 38 79 121 202 39 80 122 303 40 81 123
404 41 82 124 1717 1 42 83
1709 9 2 9 14 9 84 1701 9 11 9 52 9 93
1702 9 20 9 61 9 102 1703 9 29 9 7010 111

ZZZZZZ

32276 32643 31643 32529 31042 32464 30415 32430 30315 32428 32671 32751
32551 32739 32220 32703 31600 32685 30875 32721 29937 32888 28579 33309
27186 33992 25608 35118 24093 36609 22813 38212 32876 38515 22689 37634
22646 25132 32903 25005 32768 32768 32768 32766 32835 32764 33147 32752
33879 32718 34796 32670 35707 32618 36452 32576 37277 32525 27338 32789
27651 32977 28139 33208 29013 33466 30106 33539 31162 33393 32026 33111

444444

80 39114753 32260371 3000000 1
32515 32889 32720 32788 32686 32677 32495 32459 32093 31946 31484 31061
32768 32768 32796 32768 33394 32766 34397 32750 35249 32750 36038 32738
30709 29671 30030 28166 29352 26282 28965 24920 28762 24099 32767 32775
32656 32640 32359 32311 31934 31814 31299 31058 30593 30187 29866 29260
29570 28866 32734 32765 32480 32687 31759 32480 30491 32091 28789 31512
25865 30359 24031 29521 22186 28587 20329 27560 32739 32755 32542 32699
31915 32521 31114 32363 30505 32281 30163 32258 32617 32748 32311 32717
31633 32689 30534 32771 29290 33066 27387 33890 25578 35147 24032 36653
22585 38504 32555 37943 22085 37135 22100 24356 32408 24097

SSSSSS

4 2 7 197
101 66 131 194 202 67 132 195 303 68 133 196
404 69 134 197 1717 1 70 135 3939 11 80 144
1709 9 2 9 71 8 136 1739 8 12 8 81 8 145
1701 9 26 9 89 9 153 1702 7 35 7 98 8 162
1703 9 42 9 105 9 170 1704 6 51 6 114 6 179
1705 9 5711 120 9 185

ZZZZZZ

36915 32724 37670 32712 38708 32678 39629 32658 39643 32658 32702 32688
32349 32312 31686 31674 31201 31250 30476 30710 30000 30402 29323 30028
28805 29796 28362 29628 32764 32760 32410 32548 31629 32000 31164 31634
30418 30980 29959 30534 29375 29920 28718 29146 28385 28718 32723 32776
32264 32818 31603 32848 30856 32862 28687 32702 27157 32406 26391 32198
25477 31896 24734 31598 36222 39060 24079 37892 24085 24120 36247 24146
32768 32768 32785 32768 32793 32768 33573 32762 34551 32756 35509 32746
36403 32738 37185 32726 38251 32696 39795 32654 32759 32740 32474 32368
31886 31646 31193 30904 30555 30314 29862 29764 29190 29316 28863 29138
28049 28770 32748 32748 32225 32420 31449 31860 30609 31158 29837 30414

444444

86 53512463 19260371 6040000 4
29323 29866 28491 28868 32759 32762 32165 32774 31008 32718 29846 32578
32768 32768 32768 32768 32865 32767 33389 32770 34452 32789 36112 32789
28712 32362 27936 32154 27174 31912 26483 31654 25594 31262 36570 38294
38100 32769 40380 32728 41771 32679 43362 32628 29961 26065 30295 27427
30821 28970 31356 30236 31872 31269 32301 32027 32557 32438 32706 32677
32628 32684 32227 32456 31583 32154 30378 31746 28954 31558 27389 31695
24469 37146 24482 23370 36577 23402 32768 32768 32782 32768 32797 32768
33586 32740 34589 32704 35336 32676 36149 32642 36940 32604 38438 32528
25934 32139 24478 32892 23278 33772 32683 32768 32400 32775 31604 32792
30232 32854 28575 32960 26700 33125 24798 33347 23244 33573 22287 33728
32606 32794 32346 32826 31711 32899 30684 32967 28471 32852 26623 32491
32756 32754 32476 32388 31667 31394 30916 30602 30211 29938 29200 29144
24968 31944 23814 31447 23031 31038 32685 32789 32348 32893 31506 33148
30297 33459 28548 33808 26931 34029 25508 34154 24515 34198 23803 34204
28441 28660 27682 28270 26961 27976 32720 32750 32043 32340 31199 31746
30260 30944 29548 30256 29035 29726 28329 28882 27680 28022 27026 27060
32691 32794 32386 32902 31844 33078 30275 33402 28915 33482 27200 33309
25789 32916 25545 32825 25259 32703 40412 37680 30276 36763 30220 24172
32713 32768 31718 32726 30626 32604 29669 32448 28706 32238 28008 32038
40476 24053 32768 32768 32768 32767 32888 32771 33464 32770 35428 32786
27227 31776 26495 31504 25383 31002 36231 37482 24035 36700 23917 22846
36084 22468

SSSSSS

4 1 4 121
101 39 78 118 202 40 79 119 303 41 80 120
404 42 81 121 1717 1 43 82
170910 2 9 44 8 83 1701 9 12 9 53 9 91
1702 9 21 7 62 9 100 1703 9 30 9 69 9 109

ZZZZZZ

37846 32772 40902 32710 42788 32650 44659 32564 45618 32511 30109 25834
30336 26908 30716 28249 31265 29761 31861 31064 32321 31987 32582 32455

Figure C.2 - Sample output from SORT

444444
 80 52112922 12150471 6400000 . 1
 4 2 6 158
 101 50 101 155 202 51 102 156 303 52 103 157
 404 53 104 158 1717 1 54 105 5555 9 62 114
 1709 7 2 7 55 8 106 1701 8 10 7 6310 115
 1702 7 18 8 70 8 125 1703 8 25 7 78 7 133
 1704 8 33 8 85 7 140 1705 9 41 8 93 8 147

32768 32768 32768 32768 32904 32764 33132 32762 33504 32750 33858 32736
 34850 32690 35741 32650 32796 32668 32768 32768 32675 32668 32524 32502
 32155 32088 31336 31146 30461 30092 29570 28964 28850 28000 32768 32768
 32627 32682 32092 32390 31502 32106 29730 31538 28168 31380 26714 31508
 32759 32768 32603 32776 31980 32822 30650 32910 28267 33028 25330 33128
 23384 33156 22759 33160 32764 32766 32655 32840 32252 33090 31720 33426
 30888 33988 29653 34872 28099 36078 25938 37948 32756 32776 32670 32884
 32266 33362 31924 33744 31068 34650 29843 35790 28436 36898 26914 37874
 25550 38564 29691 41752 17610 40692 17586 26916 29676 26846 32768 32768
 32768 32768 32768 32768 33014 32764 33137 32758 33326 32748 33565 32740
 33757 32728 32793 32674 32768 32768 32613 32604 32463 32428 32057 31968
 31331 31108 30471 30048 29357 28602 32748 32752 32634 32676 32154 32356
 31488 31978 30688 31610 29086 31144 28141 31032 27053 31060 32743 32766
 32366 32798 31636 32846 30050 32944 28496 33028 25165 33146 23144 33184
 32768 32768 32631 32850 32446 32966 32212 33118 31971 33270 30901 33984
 29650 34884 28554 35728 32762 32776 32646 32912 32261 33368 31725 33970
 30493 35216 29090 36428 26866 37922 26285 38236 30167 40936 18031 39962
 17981 26154 30109 26004 32768 32768 32768 32768 32768 32768 32818 32768
 32961 32764 33047 32760 33367 32754 33690 32742 33874 32736 32800 32690
 32768 32768 32661 32652 32520 32490 32189 32114 31780 31632 31382 31154
 31006 30690 30767 30388 30175 29638 29628 28916 32768 32768 32596 32634
 32493 32528 31975 32192 31229 31738 30102 31196 27707 30574 25374 30690
 32742 32766 32378 32786 31792 32824 30127 32902 27859 32982 24011 33046
 22730 33046 32768 32768 32628 32854 32204 33110 31702 33428 29639 34832
 28091 36012 26842 37046 32768 32768 32644 32908 32487 33098 32160 33472
 31515 34168 30167 35458 28975 36440 27830 37232 35074 40110 17483 39006
 17474 25134 35044 25092

222222
 444444
 80 52112861 12150471 6440000 . 1
 4 3 6 159
 101 52 104 156 202 53 105 157 303 54 106 158
 404 55 107 159 1717 1 55 108 5555 10 65 115
 3939 11 66 116
 1709 8 2 8 57 6 109 1739 7 12 8 67 6 117
 170110 19 7 75 9 123 1702 7 29 7 82 8 132
 1703 9 36 9 89 8 140 1704 7 45 6 98 8 148

32768 32768 32768 32768 32768 32768 32893 32764 33145 32754 33502 32736
 33830 32718 34162 32702 34604 32678 32843 32850 27893 30128 30464 32430
 30610 32500 31116 32656 31560 32758 32004 32798 32423 32792 32685 32772
 32768 32768 32616 32716 32426 32636 31945 32462 30836 32082 29782 31752
 28667 31438 27633 31170 26588 30938 24031 30488 32745 32762 32434 32748
 31747 32722 30531 32772 28046 33244 24734 34668 23422 35492 32758 32770
 32368 32804 31776 32844 30797 32912 29737 32976 27578 33094 23313 33228
 22305 33242 19370 33244 32761 32772 32514 32888 32192 33040 31211 33474
 29541 34098 28718 34354 27770 34604 29175 40906 17078 39912 17040 26138
 29144 26000 32768 32768 32768 32768 32768 32768 33022 32760 33269 32748
 33636 32730 34054 32708 34498 32682 34817 32668 32847 32856 27475 30888
 27470 33892 30555 32850 30773 32886 31330 32930 31726 32928 32197 32882
 32569 32814 32732 32784 32768 32768 32632 32724 32203 32568 30966 32154
 29851 31824 28478 31460 25241 30804 32768 32768 32570 32762 32047 32752
 31240 32782 30455 32864 28798 33176 27276 33658 32766 32774 32376 32796
 31841 32830 31410 32860 30375 32926 28293 33030 27146 33078 25123 33140
 20159 33192 32768 32768 32645 32820 32159 33046 31420 33348 30062 33844
 28493 34296 29633 40070 17496 39120 17447 25304 29571 25142 32768 32768
 32768 32768 32866 32768 33108 32768 33419 32760 33678 32752 34054 32740
 32841 32862 27851 31576 30477 33086 30754 33130 31154 33130 31589 33092
 32229 32956 32745 32784 32768 32768 32635 32724 32389 32636 31866 32450
 31071 32180 29851 31814 28425 31430 26842 31074 25542 30828 32742 32764
 32375 32764 31789 32772 31124 32822 29078 33182 25849 34454 24519 35232
 24519 35232 32768 32768 32396 32786 32034 32806 31239 32844 29381 32902
 28135 32934 24846 32974 23292 32966 32768 32768 32625 32824 32254 32966
 31716 33170 30250 33642 29892 33888 27967 34158 27967 34158 29135 39224
 16947 38074 16951 24196 29118 24202

222222

40 words of sector zero of working storage are not used to hold a card image, but are used to hold three numbers which inform SORT how much of the space in working storage is taken up by the programme. SORT can then transfer the data onto the disk starting with the first available space following the programme.

SORT will output the data back onto the same disk from which it read the data, or onto another disk, whichever is required. Usually two separate disks are used, partly because the input disk has so much data on it that there is no room for the sorted output, and partly because SORT executes more rapidly when it is transferring data from one disk to another rather than from one part of a disk to another part of the same disk.

(c) Operation

The programme SORT provides several options which are selected by different settings of the console entry switches. A printed request is given to inform the operator when and how these settings should be made. A brief description of the different options is given here. Options (i) (ii) and (iii) are selected during the initial stages of the programme, and the others may be changed at any time during execution.

(i) Logical drive number of the output disk. If the data are to be output onto the working storage of the master cartridge (the first one specified on the JOB record), then no switches need to be set. If the data are to be output onto the second disk specified on the JOB record, then switch 15 should be set.

(ii) Choice of printer. If a list of the events and their heading information is required this will be given on the console printer if switch 1 is set. Alternatively, if the programme is being run on the terminal computer for NUMAC which has a 1403 printer, then SORT can take advantage of this faster printer if switch 0 is set.

(iii) Starting sector. Usually SORT is requested to start scanning the data from the beginning, in which case no switches need to be set. It is occasionally useful however, to request SORT to start scanning part way through the data. The number of sectors to be skipped is indicated by using the switches as a 16-bit binary counter (bit 15 is the least significant bit).

(iv) Selection of machines. Console entry switches 1 to 6 are used to inform SORT which data are to be prepared for input to NUMAC. Only data from those machines whose corresponding bits are on will be selected.

(v) Selection of events. Usually each complete event read from DATA1 is transferred to the output disk, subject to the selection by machine number as indicated in (iv) above. Sometimes however, it is required to output only selected events, or to suppress the output of selected events. This is achieved by the use of switch 0 which should be set for each event whose output is to be suppressed. The heading information for the event still appears on the listing of the events, but it is preceded by "***" to indicate that the data for this event were not transferred to the output disk.

(vi) Test. The final option gives the operator the chance to examine the heading information of an event before deciding whether or not to suppress the output. When bit 15 is set, SORT prints the heading information for the event and then waits. If bit 0 is on when execution is resumed the output is suppressed. If bit 0 is off, the data are output onto disk and the heading information printed again to record that this has been done.

SORT begins by selecting events for machine 1 if that switch is on, and will continue to scan for machine 1 events until it reaches the end of the data, unless switch 1 is turned off, in which case no more events for machine 1 will be selected. Before proceeding with machine 2 events, SORT prints a message which informs the operator how many events for machine 1 have been transferred to the output disk.

This procedure is repeated for each of the six machines.

APPENDIX DConsole printer messages

The console printer is used to provide a log of the measuring and to give the operators coded error messages. It may also be used to provide an operator with more information about what is required of her than can be given on the lamp display.

Each message is preceded by the number of the machine to which it refers, and is tabulated across the page according to its machine number. All messages are kept as brief as possible to minimise printing time.

As part of the response to an ACCEPT DATA interrupt which is used to request the start of a new event, the event number which is read in from the decade edge switches is printed.

e.g. M1 36507562

If the coordinate measurements of a track fail to satisfy the curve fitting routine, the extent of the failure is indicated by a printed message.

e.g. M1 aaaa bbbb

where aaaa is the distance in fringes of the apex from the fitted circle, and bbbb is the standard deviation in fringes of the fitted circle.

Similarly, if a track measurement fails to satisfy the sagitta test, this is indicated by a printed message.

e.g. M1 9999 cccc

where cccc is the length of the sagitta in fringes.

If an operator wishes for more information than is given on her lamp display, a CHECK interrupt request will cause the label of the measurement

currently being requested to be printed.

e.g. M1 aa
or M1 ab nn

where aa is the label of a point, ab is the label of a track, and nn is the number of points already recorded for that track. This message is also printed when an erasure is made, when an apex check is unacceptable, when an apex setting value is unacceptable, and when an END TRACK request is given with fewer than 5 prints recorded on the current track.

The operator may force the programme to accept a track measurement which has failed, by using the ACCEPT interrupt request. A printed message records that this has been done.

e.g. M1 ab AXPT

where ab is the label of the track which has been accepted.

Finally, there are a number of coded error messages, and one warning message, which are given to inform the operator when she makes an invalid interrupt request. These are shown in Table D.1.

Table D.1

Coded warning and error messages

Message	Error condition
WARN 01	Request to sign off during the measurement of an event
ERR 00	Contact bounce on the EXECUTE interrupt button
ERR 01	Request to start a new event before the operator has signed on
ERR 02	Request to sign on during the measurement of an event
ERR 03	Request to sign on before the previous operator has signed off
ERR 06	Invalid topology coding
ERR 07	Topology specifies an event which requires more than 8 point labels
ERR 08	Topology specifies an event which requires more than 10 track labels
ERR 09	Request to enter more than 15 points for one track
ERR 10	END TRACK interrupt raised after track failure due to fewer than 5 points, or too short a sagitta
ERR 11	END TRACK interrupt raised after curve fit failure
ERR 13	ACCEPT interrupt requested without END TRACK being tried first
ERR 20	Interrupt request other than ACCEPT DATA, END or START with no event in process of being measured
ERR 21	No switches set in illuminated switch panel after an EXECUTE interrupt has been raised
ERR 31	CHECK/READ POINT interrupt requested before the view has been started

APPENDIX EList of Abbreviations

Abbreviation	Meaning
BOSC	Conditional Branch or skip (turn off interrupt level)
BSI	Branch and store instruction address register
CES	Console entry switches
CPU	Central processing unit
DIPW	Digital input word
DOPW	Digital output word
DSW	Device status word
EBCDIC	Extended binary coded decimal interchange code
ILSW	Interrupt level status word
I/O	Input/output
IOCC	Input/output control command
NUMAC	Northumbrian Universities Multiple Access Computer
OEM	Original equipment of the manufacturer
SAC	Storage access channel
SLA	Shift left accumulator
WHDSW	Westinghouse device status word
XIO	Execute input/output

REFERENCES

1. LORD, D., 1967, CERN Report DD/DA/67/8.
2. MARCH, P. V., 1968, Private communication.
3. DONALD, R., 1968, Private communication.
4. MacEWAN, J. F., 1968, Private communication.
5. LAWRENCE, D. E., 1968, Private communication.
6. CANDLIN, D., 1969, Private communication.
7. WERBROUCK, A., 1970, Private communication.

ACKNOWLEDGEMENTS

The author wishes to take this opportunity of thanking Professor G. D. Rochester, F.R.S. for his interest and encouragement and for providing the facilities which enabled the author to develop the on-line measuring system.

The author is greatly indebted to her supervisor, Dr. J. V. Major, for his constant guidance and for many useful suggestions.

All the members of the High Energy Nuclear Physics Group are thanked for their cooperation, and the Durham University Computer Unit is thanked for the use of the NUMAC terminal computer for the preparation and testing of many of the programmes.

Finally, the author would like to thank Westinghouse Brake and Signal Company for permission to copy one of their drawings (Figure 3.1).

