SLAC-PUB-752
May 1970
(EXPI)

STATUS OF HUMMINGBIRD FILM DIGITIZERS*

John L. Brown

Stanford Linear Accelerator Center
Stanford University, Stanford, California 94305

In this report I would like to review briefly the status of CRT film digitizers at SLAC. I will start with a short description of the hardware, and then summarize our experience to date on three different experiments. I will omit any discussion of the Spiral Reader, although it is also a part of SLAC's automatic data analysis effort.

I. BRIEF DESCRIPTION OF HARDWARE AND COMPUTER CONFIGURATION

A. 360/91 Computer

Our CRT film digitizers are connected online to an IBM 360/91 multiprogrammed computer. The overall computer configuration is shown in Fig. 1. It is obviously a fairly complex system; most of it need not concern us, however, save for the 2250 display scope (which is used for online interaction), two disk drives (used for storage of programs and data) and a high-speed selector channel to which our hardware is connected via a 2701 parallel data adapter.

In its current configuration the 91 is a fairly powerful machine, with the throughput of roughly 2 CDC 6600's. It is a multiprogrammed machine, usually processing half a dozen jobs at once: a mixture of batch and express jobs, terminal programs and one or more "online" programs. Typically up to a thousand separate

(Presented at Second International Colloquium on PEPR,
Massachusetts Inst. of Tech., May 5-7, 1970)

*Work supported by the U. S. Atomic Energy Commission.

jobs are processed each day. The facility is open for normal operations from about 10:00 am to 3:30 am Monday morning to early Saturday morning.

A recent development in the operation of the facility, whose impact upon automatic data analysis has not yet been evaluated, is the rationing of computer use. For a number of months users have had an accounting of their running on the 91. Use is measured in so-called "computer units", determined by a fairly elaborate algorithm based on one's use of core, CPU cycles, I/O accesses, and so forth. One month ago users were restricted to using only a predetermined number of computer units per quarter. To the extent that demand exceeds supply, this is going to provide quite an impetus toward efficient programming.

B. Hummingbird II

The Hummingbird II is by now rather an elderly CRT flying spot digitizer. It uses a 7" Ferranti 7/29AO cathode ray tube to generate a 65 × 105 mm raster using essentially 1:1 optics. The deflection and focusing coils are manufactured by Celco, while the analog electronics are homebuilt. The digitial logic is based on DEC cards, although as we shall see later, it is being converted to IC logic. The film transport is rudimentary, using a stepping motor to drive 70 mm perforated single strip film.

The raster is composed of 4096 least counts. The least count on the film in the X direction (along the line) is $\sim 4.7\,\mu$, while in the Y direction it is 25μ. Only static pincushion correction is used, so the raster is fairly noticeably distorted. At the start of each line a y-coordinate is read out to the channel; if the spot crosses a black mark on the film, the x coordinate of the center is read out. The center coordinate is determined by delaying the PM pulse and detecting the crossing point, provided the signal exceeds a certain threshold. No pulse-height or width information is supplied. Thus the output to the computer consists of a string of two-byte words as follows:

$$\ldots\ O, Y_j, X_{j,1}, X_{j,2} \ldots X_{j,n}, O, Y_{j+1}, X_{j+1,1} \ldots X_{j+1,n}, O, Y_{j+2}, \ldots$$

The zeroes are used to identify the following half word as a y coordinate.

The scanner is capable of executing a fairly limited repertoire of commands from the computer:

(1) raster-scan an area beginning at Y_i and ending at Y_f, with the PM gated on between X_i and X_f;

(2) select a line density of every line, every other line, or every fourth line;

(3) set the PM threshold to one of sixteen values;

(4) advance (or back up) the film up to 20 frames, in increments of 1% of a full frame advance.

C. Hummingbird III and TV Display

Hummingbird III is very similar in overall design to HB II. It uses a 9" Litton L-4192 pentode CRT with a P24 phosphor to generate a spot which is imaged with custom made 1:1 Zeiss optics onto a film platen — field lens — photomultiplier unit. The film drive is designed to handle 3-strip perforated 35 mm film. It contains a pneumatically driven carriage assembly which moves up and down in a vertical plane to position the appropriate view (or a calibration pattern) over the fixed vacuum platen. Celco focusing and deflection coils are used, driven by electronics based on Beta Instrument Company circuits. Digital logic is made from IC's. The logic design is such that it will eventually link both HB's to the 91 channel.

HB3 was designed for use with bubble chamber or streamer chamber film, which can produce upwards of 50K digitizings per frame. Since the IBM 2250 display scope can only hold about 1300 points in its buffer, a different type of display would be desirable. We have constructed a digital TV display using a conventional industrial TV monitor with a 512 × 512 raster. The picture is stored on a fixed-head disk (manufactured by Data-Disc) which refreshes the interlaced image every thirtieth of a second, as a normal TV set does. A lightpen is attached which stores its recorded data on a separate set of tracks. Also available are a separate set of tracks for display of points with enhanced brightness. A program function keyboard is also included. This TV scope is capable of displaying 100K points without flicker. There is a fairly high software overhead in converting a FORTRAN array of points, vectors, or characters into the appropriate bit string for storage on the disk.

II. μ-p SPARK CHAMBER EXPERIMENT

Our first experiment, completed last December, was a spark chamber experiment designed to study μ-p elastic and inelastic scattering, to see if the muon and electron exhibited any differences in this respect. A sketch of the experimental layout is shown in Fig. 2. The muons, after scattering from the target, passed through two spark chambers, a 54" momentum-analyzing magnet, two more conventional chambers, and then four more chambers interspersed with absorber to distinguish muons from pions. 90° stereo views of all eight chambers were taken. In addition there was a proton recoil chamber mounted underneath the target which we did not use in our analysis. There were a total of 10 fiducial marks (each in the shape of a V), and a data box containing a BCD representation of the roll and frame number.

The overall program design changed somewhat as we moved from the interim 360/75 to our current 360/91 computer. In the final version the program was a single package occupying 300K bytes of core. The program drove the scanner in a buffered manner, i.e., while the current frame was being processed, the next frame was being digitized. Since only one frame in three contained a real event, the film had been rapidly prescanned by human operators. The first processing of a frame (typically containing 4000 digitizings or "hits") consisted of stringing "hits" together into small straight line segments (called "blobs"). The 400 or so "blobs" which resulted represented fragments of tracks, fiducial marks, data box bits, scratches, etc. The fiducial marks and data box were then found and checks made for fiducial separation, data box parity errors, etc. Next the remaining "blobs" were sorted into the expected chamber locations and connected where possible to form "segments," i.e., images of a single track in a particular view of a particular chamber. Then segments were joined to form complete trajectories, after making allowances for displacements and rotations caused by the spark chamber optical system.

If the program could not come up with a single unambiguous "goldplated" event, matching the description on the input scan card, then the program halted for operator intervention at the 2250 display scope; this was the case on 75% of the frames, so the data analysis was scarcely "automatic." The operator could link "blobs" into "segments" or "segments" into tracks using the light pen. With this manual intervention, the program could process events at the rate of 60-100 per hour.

In the course of processing some 125K frames on this experiment from March 1968 to December 1969, we came by several hard-learned lessons.

a. The program design, which started of course before the experiment was run, naively assumed that the pictures would be "perfect." We didn't allow for the fact that half the fiducial marks (made from electroluminescent strips) would burn out in the course of the experiment, as would a number of the bits in the data box. We didn't make allowances for "ghost" tracks caused by reflections in the Lucite walls of the chambers, and we had trouble as well with variation of spark intensity as a function of the number of tracks in a chamber. Occasional low chamber efficiency also caused us to miss tracks.

Some of these problems can be circumvented. For example, you can cover the Lucite walls with black paper to cut down reflections, but you have to think of it before you take the pictures, not afterwards. Some problems can be dealt with in the scanner hardware (e.g., better track-center circuits which work over a wider range of image contrast), and some can be overcome in software (e.g., better track-finding algorithms that don't assume effectively that every gap in every chamber will fire). It does seem to be a fact of life that you don't learn these lessons from reading about them, but only by having them happen to you.

b. A second problem which we generated for ourselves was to try to cover up failures in the scanner hardware with software "fixes." The particular problem we had was quite complex and difficult to explain, and would probably not be of general interest. It had to do with the way in which we calibrate the scanner, which is to scan a pattern of 54 crosses whose center positions are accurately known in a rectilinear coordinate system. A fifth-degree polynomial is used to transform HB coordinates into true film coordinates, with the transformation coefficients being determined from a scan of the cross pattern. What happened to us was that, due to gradual misalignment, the spot size in the corners of the raster got so large that the crosses in that area weren't properly digitized. The calibration routines then omitted them from the fitting procedure used to find the transformation coefficients. If the various distortions (such as pincushion) are large, however, (and they are since we don't use dynamic distortion corrections) then the transformation coefficients are very sensitive to the presence or absence of these corner crosses. In our output this would manifest itself as small shifts in angles, for example, compared to hand measurements, and these shifts would vary with time, depending upon how many crosses in the calibration pattern had been well digitized. The proper solution, which we finally adopted was to stop and realign and tune up the scanner, rather than try futilely to remedy the problem with software changes.

c. A third lesson we learned, as has everyone else before us, is the importance of having physics analysis programs ready before vast numbers of measurements are accumulated. In our case this meant having a well understood and debugged geometry program for fitting an overconstrained trajectory through the magnet, based on track measurements in the various chambers. Such a program wasn't available for us until we had measured the majority of the film, at which time it uncovered the problems referred to in the previous section.

As a result of all these problems, our first experiment was a mixed success. Production went in fits and starts as problems were uncovered. In the end it turned out that most (80-85%) of the events in the experiment were not muons scattered from the target, and HB measurements were trusted when they indicated this. If the event appeared to come from the target, however, it was remeasured by hand, since that was considerably more straightforward than trying to understand the milliradian systematic errors present in HB output.

III. COSMIC RAY SPARK CHAMBER EXPERIMENT

A second experiment in which we are currently involved is a cosmic ray spark chamber experiment. This is a collaborative effort between SLAC and LRL. It is designed to measure the momentum and angular spectrum of cosmic ray muons at sea level, and in particular to check the zenith angle distribution of the highest energy muons. About 1.6×10^6 pictures have been taken. A rapid hand scan is being done to pick out the highest energy muons for subsequent hand measurement with the greatest possible precision. The other 98% of the data is to be analyzed on the Hummingbird, where a slightly lower precision is acceptable.

The experimental layout is shown in Fig. 3. The apparatus is by and large the same as was used in the μ-p experiment, but slightly rearranged. There are three chambers for determining the trajectory of the muon before the momentum analyzing magnet, and three after. There is also one chamber inside the magnet. All but one of the chambers have a 90° stereo view. There are 20 V-shaped fiducials and a BCD roll-frame data box. Counter information was recorded at the time of the experiment by a PDP-8 and this data is available for merging with Hummingbird output.

The overall program design is similar to the μ-p experiment. The main difference is that the track-finding algorithms are more global, and don't depend so much on precisely what is happening in a given chamber. There is also a second pass feature in the program whereby if an event isn't found using the "blobs" then one can go back to the original digitizings to see if an event can be found.

The film is much "cleaner" than the μ-p film was, largely as a result of the lessons learned in the latter experiment. (The duty cycle is also better: 100% instead of 0.1%). Consequently our track-finding efficiencies are better. Currently we correctly resolve about 85-90% of the frames. About two-thirds of the remainder have no real events in them at all, while the other third (about 3-5% of the total) contain events of varying degrees of complexity (e.g., showers). The goal in this experiment is to do without manual intervention for event finding, and we are fairly close to achieving this. The processing rate is about 700 frames per hour, limited essentially by the rate at which the Hummingbird can scan and move film. A factor which may limit our overall production rate, however, is the rationing of computer units referred to earlier. Our allocation is such that we may be limited to about 8 hours a day of production rather than the potential 16. An interesting sidelight to this accounting and budgeting problem is that the computer costs for a single frame of cosmic ray film are currently about eight cents.

The biggest single problem remaining in the cosmic ray experiment is the question of the unresolved events. In an experiment of this magnitude and potential statistical precision, the fraction of unresolved events should ideally be about 2%, instead of the current 10-15%. Since the film has not been completely prescanned, the problem is as much one of deciding there is no event as of finding one which is there. It is not clear what strategies will be used to solve this problem.

The second problem is to keep a close watch on potential small systematic distortions caused by the Hummingbird hardware. Since we do have the most important spatial reconstruction programs online, we can monitor our accuracy much better than we could in the μ-p experiment.

IV. STREAMER CHAMBER K_2^o DECAY EXPERIMENT

This experiment, on which we are starting some shakedown runs, is our most ambitious to date. This experiment, a collaboration between SLAC and BNL, is designed to study leptonic K_2^o decays using a streamer chamber as detector. Absorbing plates are put in the chamber to help separate pions, muons and electrons. A sample picture is shown in Fig. 4. This film (3 strip 35 mm perforated) will be digitized on HB III. A small amount of film has been taken on this experiment already, but the bulk of the data will not be collected until June.

The nucleus of the streamer chamber software is CERN's Minimum Guidance program. The film will be prescanned (since only one frame in 5 contains an event) and a rough vertex position recorded on a scan card. The Hummingbird will then scan all three views with a "normal" scan, and an orthogonal scan if so indicated on the scan card. A vertex-finding program then uses the rough vertex position to find a precision (100 μ on the film) vertex which serves as input to the Minimum Guidance program. The MG program then follows the decay products emanating from the vertex (into the orthogonal scan if necessary) until the tracks reach the absorbing plates. A "pseudo-vertex" is then constructed on the exit side of the plates to follow the tracks as they emerge (if they do) from the other side. The pieces of track found by the MG program in the three views are then edited, labelled, and checked for topological consistency. If it appears that an event matching the description on the scan card has been found, then the measurements will be given to SYBIL, a three-view geometry program similar in purpose to TVGP but adapted to the peculiarities of the streamer chamber.

If at any point along the chain the program experiences difficulty then manual intervention from the TV scope is called for. A sample display (used for debugging, not production) is shown in Fig. 5. The operator can then erase irrelevant digitizings, link track segments, indicate vertices, etc., using the light pen and program function keyboard.

The overall program is designed to function asynchronously, i.e., the scanner fills up disk storage with digitized frames for processing, the vertex finding routine accumulates vertices, the MG program finds tracks, and messages and pictures for display are stored, all more or less independently. The idea is to avoid a strict sequential "bucket brigade" operation, and rather to have all pieces of the program working on their own input queues. Of course eventually some one part (e.g., the scope operator) becomes a bottleneck, but this asynchronous design keeps him continuously busy.

We have so little experience to date that there is nothing but bad news to report.

A. We are experiencing a great deal of difficulty with HB3 hardware. Two critical problems are noise in the IC digital logic, and stability in the focusing and deflection circuits. The track center circuit is also undergoing considerable rework. Although HB3 is similar to HB2 in overall specifications and capabilities, the actual detail design is almost completely different, so there remains considerable debugging to be done.

B. The film we have, while better than early streamer chamber film, still is far from optimum. There is still great variation in track contrast and width, as well as large flares which can obscure considerable portions of an event. It is not clear at present whether streamer chamber film in its present state is actually amenable to automatic data analysis.

C. The list of software problems is almost endless. The total program is very large, involving some 200 subroutines occupying more than on megabyte of storage. Since we are only allowed 300K of core storage, this means a great deal of overlaying both of instructions and data. Consequently there is a great deal of channel activity between disk and core, and we are currently trying to sort out this channel traffic. While the vertex finding program works well (97%) on a small sample of events, we will undoubtedly run into problems when we try to go to a larger less selected sample. The same is true of the MG program; it finds about two-thirds of the tracks in a small sample but there still remains considerable tuning of program parameters. The editing program has been only partially checked out, and no events have yet been input to SYBIL. Thus, while there has been reasonable success with individual program components, overall system checkout has not been attempted and considerable problems can be expected.

V. FUTURE PLANS

Perhaps the most accurate statement is that we are so busy with present problems we haven't had any time to work on future plans. At the moment there are no intentions of expanding or improving our hardware capabilities in any significant way. There is a specific proposal for a spark chamber experiment on electroproduction of hadrons to take place in about a year, but no serious programming work has yet been done. A rather massive streamer chamber exposure to a high energy K^- beam is also planned for the beginning of 1971, and if we have reasonable success on the K_2 experiment we can expect a large amount of work in analyzing this next experiment. And, of course as an ex-bubble-chamber-physicist, I have a personal interest in adapting our streamer chamber program to bubble chamber experiments. Unfortunately, in my current role as a bureaucrat faced with what he considers an inadequate budget, my main problem at the moment is to figure out how to do more work than we can handle with fewer resources than we need.

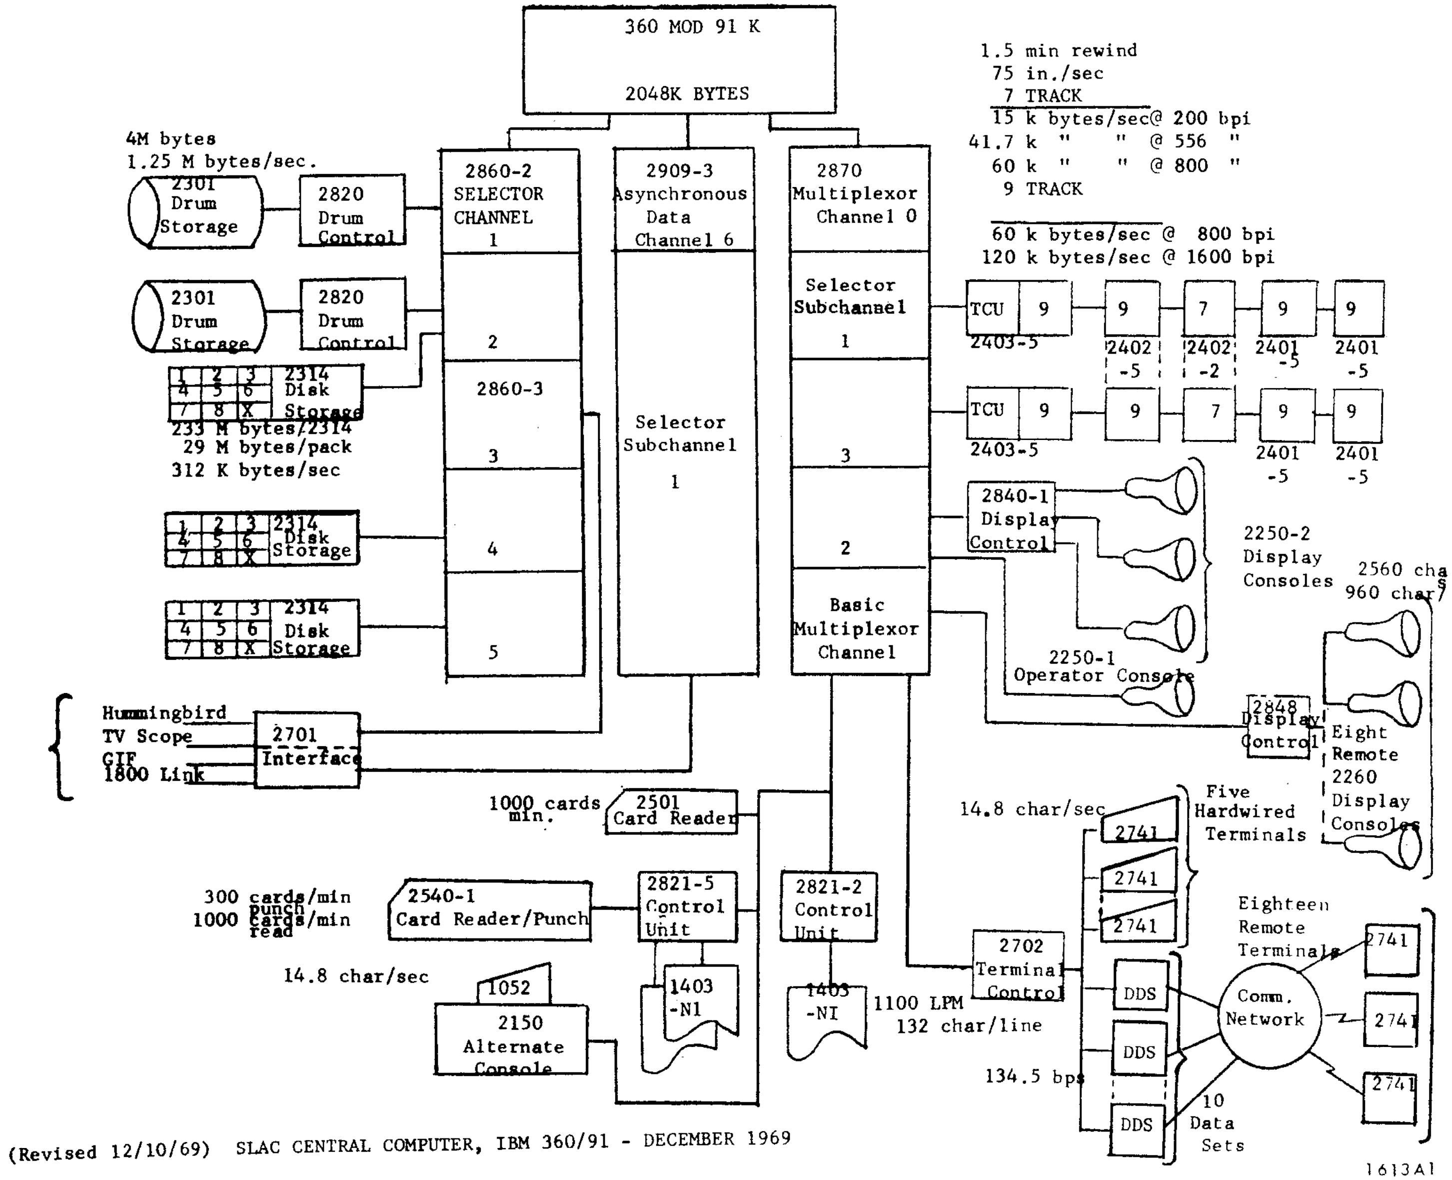
360 MOD 91 K
2048K BYTES
2860-2 SELECTOR CHANNEL
2909-3 Asynchronous Data Channel 6
2870 Multiplexor Channel 0
4M bytes
1.25 M bytes/sec.
2301 Drum Storage
2820 Drum Control
2314 Disk Storage
233 M bytes/2314
29 M bytes/pack
312 K bytes/sec
2860-3
Selector Subchannel
Basic Multiplexor Channel
1.5 min rewind
75 in./sec
7 TRACK
15 k bytes/sec@ 200 bpi
41.7 k " " @ 556 "
60 k " " @ 800 "
9 TRACK
60 k bytes/sec @ 800 bpi
120 k bytes/sec @ 1600 bpi
TCU
2403-5
2402 -5
2402 -2
2401 -5
2840-1 Display Control
2250-2 Display Consoles
2250-1 Operator Console
2560 chars
960 chars
2848 Display Control
Eight Remote 2260 Display Consoles
Hummingbird
TV Scope
GIF
1800 Link
2701 Interface
1000 cards min.
2501 Card Reader
300 cards/min punch
1000 cards/min read
2540-1 Card Reader/Punch
2821-5 Control Unit
2821-2 Control Unit
14.8 char/sec
1052
2150 Alternate Console
1403 -N1
1100 LPM
132 char/line
14.8 char/sec
2741
Five Hardwired Terminals
2702 Terminal Control
DDS
134.5 bps
10 Data Sets
Comm. Network
Eighteen Remote Terminals
(Revised 12/10/69)
SLAC CENTRAL COMPUTER, IBM 360/91 - DECEMBER 1969
1613A1

Fig. 1

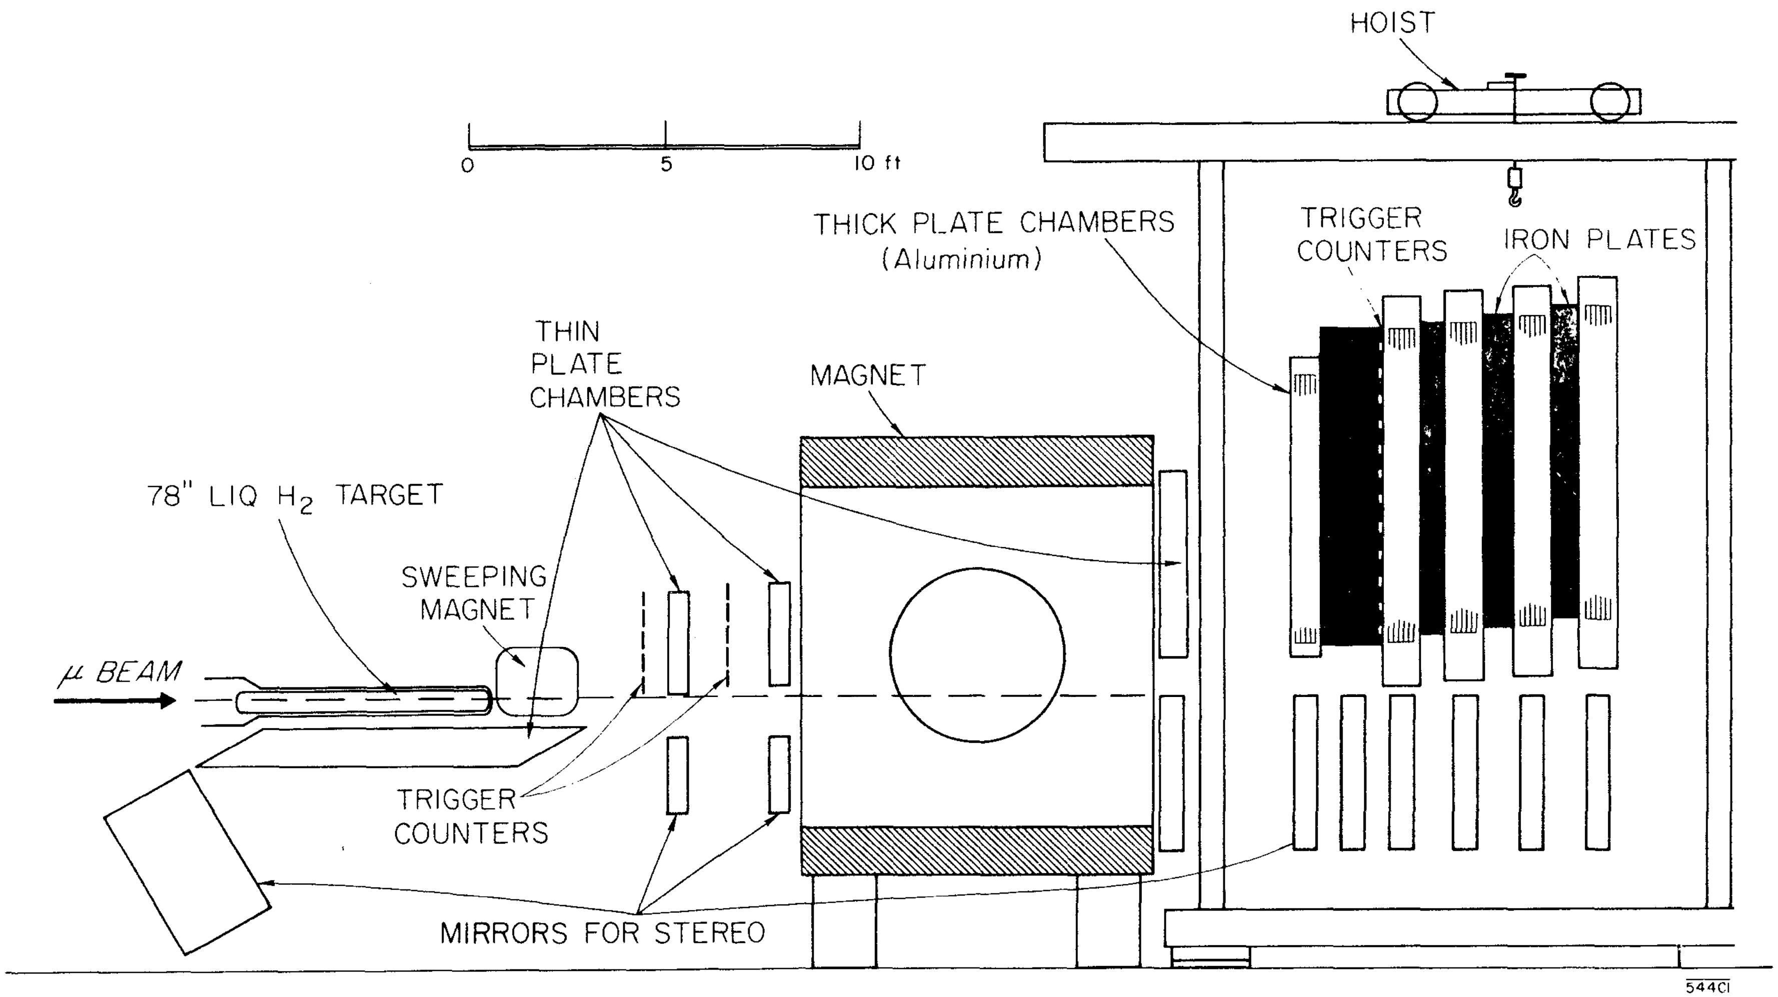
HOIST
0
5
10 ft
THICK PLATE CHAMBERS
(Aluminium)
TRIGGER
COUNTERS
IRON PLATES
THIN
PLATE
CHAMBERS
MAGNET
78" LIQ H_2 TARGET
SWEEPING
MAGNET
μ BEAM
TRIGGER
COUNTERS
MIRRORS FOR STEREO
544C1

Fig. 2


1486A2

Fig. 3

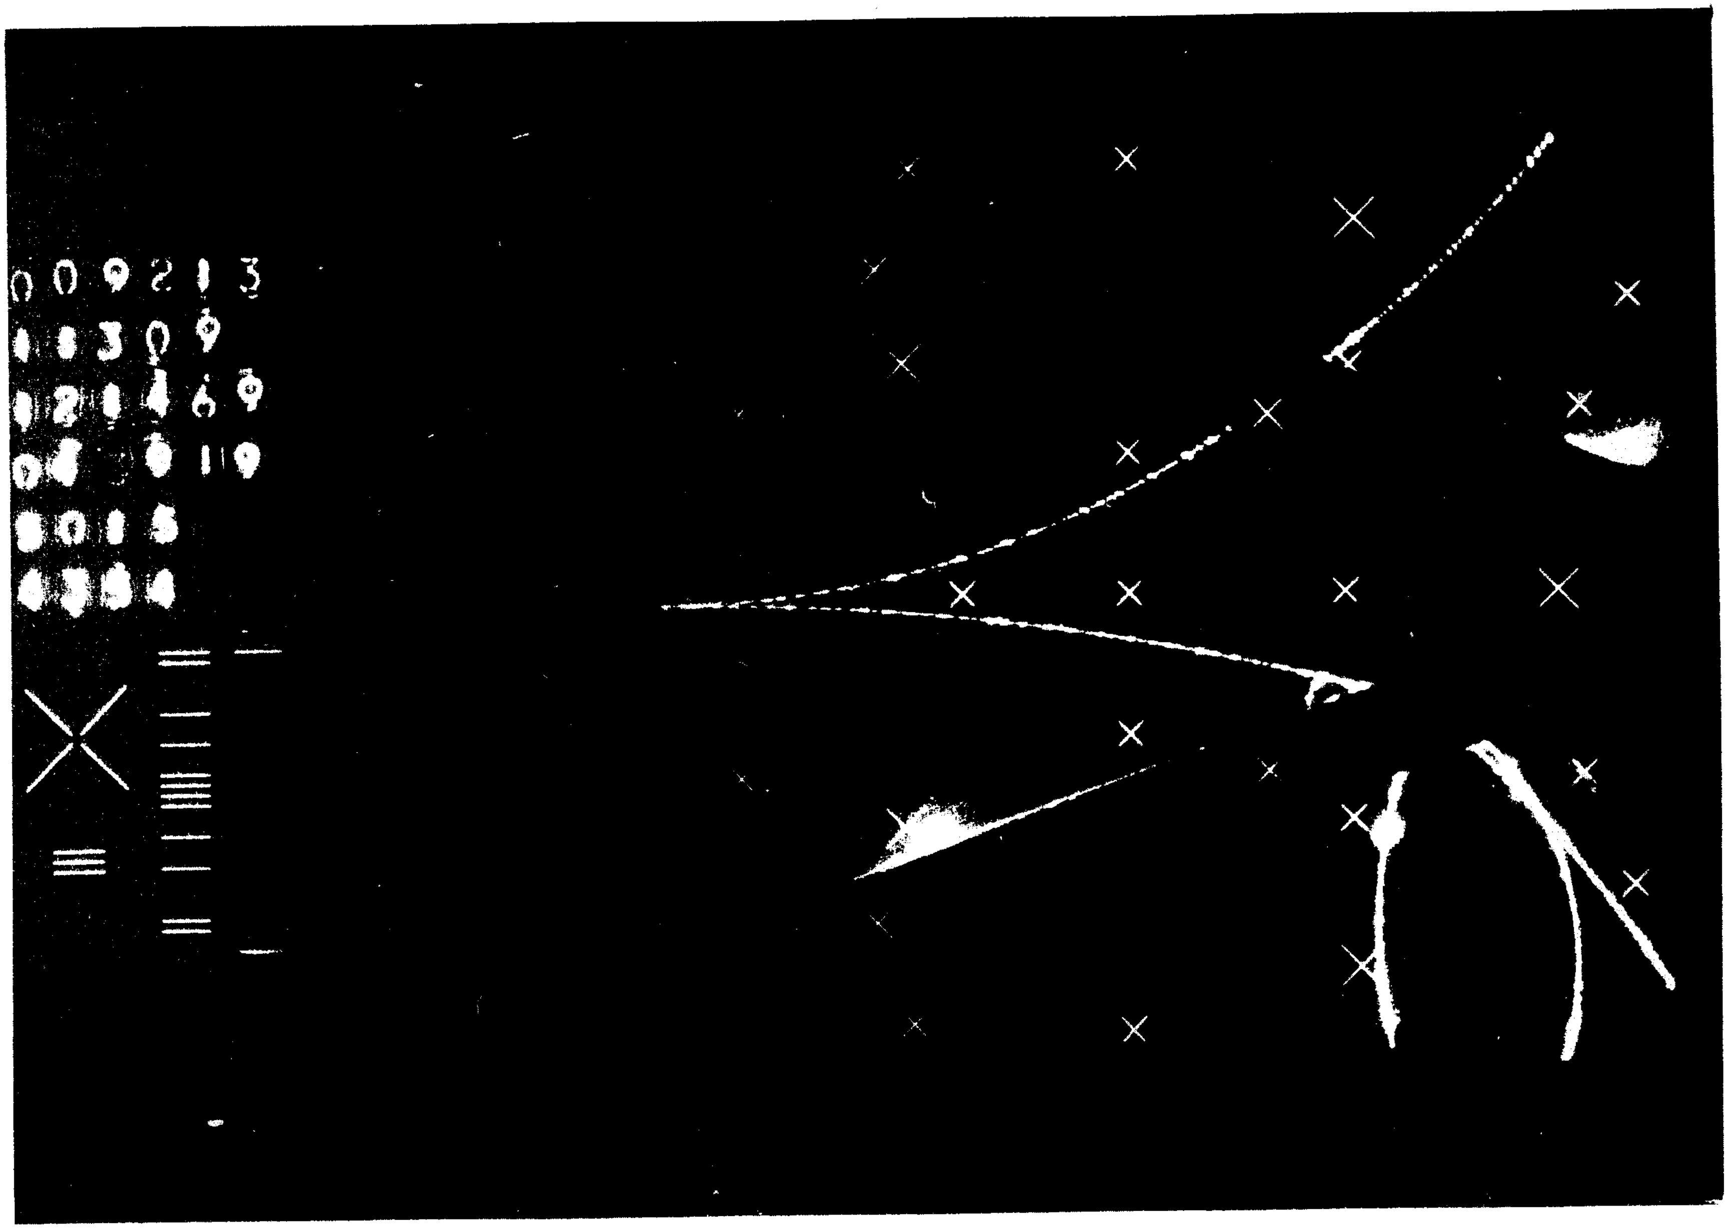

Fig. 4

1613A2

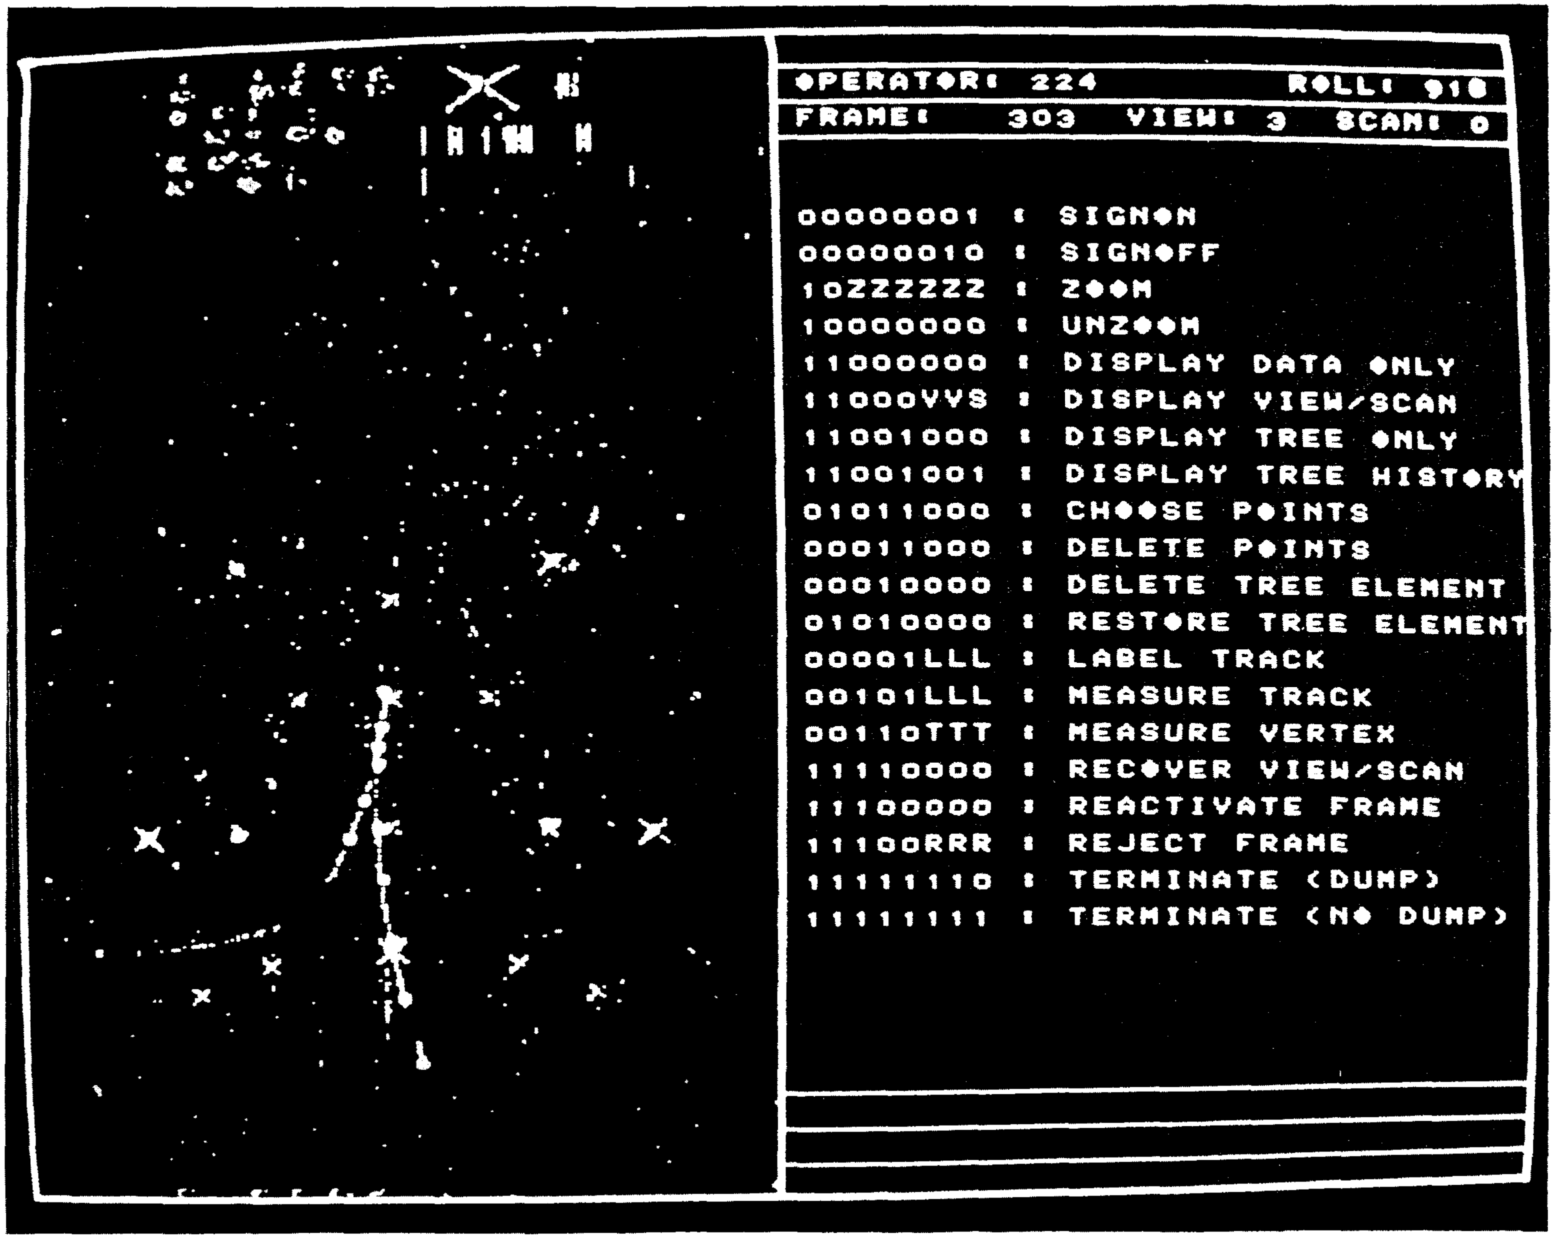
OPERATOR: 224
ROLL: 91
FRAME: 303 VIEW: 3 SCAN: 0
00000001 : SIGNON
00000010 : SIGNOFF
10ZZZZZZ : ZOOM
10000000 : UNZOOM
11000000 : DISPLAY DATA ONLY
11000VVS : DISPLAY VIEW/SCAN
11001000 : DISPLAY TREE ONLY
11001001 : DISPLAY TREE HISTORY
01011000 : CHOOSE POINTS
00011000 : DELETE POINTS
00010000 : DELETE TREE ELEMENT
01010000 : RESTORE TREE ELEMENT
00001LLL : LABEL TRACK
00101LLL : MEASURE TRACK
00110TTT : MEASURE VERTEX
11110000 : RECOVER VIEW/SCAN
11100000 : REACTIVATE FRAME
11100RRR : REJECT FRAME
11111110 : TERMINATE (DUMP)
11111111 : TERMINATE (NO DUMP)

1613A3

Fig. 5