Title
THE DAPR PRODUCTION SYSTEM

Authors
Hall, Dennis
Pinchak, Joan
White, Howard S.

Publication Date
1965-10-01

UCRL-16510

University of California

Ernest O. Lawrence Radiation Laboratory

THE DAPR PRODUCTION SYSTEM

Conf. on Programming for Flying
Spot Devices, Columbia University

UCRL-16510

UNIVERSITY OF CALIFORNIA
Lawrence Radiation Laboratory
Berkeley, California

AEC Contract No. W-7405-eng-48

THE DAPR PRODUCTION SYSTEM

Dennis Hall, Joan Pinchak, and Howard S. White

October 1965

The DAPR Production System

The previous paper has described prototype Digital Automatic Pattern Recognition (DAPR) programs which have operated at Berkeley.[1] These programs operate in conjunction with the IBM 7094 II computer and the Flying Spot Digitizer, a Hough-Powell device.[2,3] This hardware complex has been in operation since July 1963 for physics production using the HAZE system of manual scanning and automatic measurement.[4]

A limited number of frames has been analyzed from each of the Hydrogen Bubble Chambers which can be measured by the FSD. These are the Berkeley 72" and 25" and the Brookhaven 80" Hydrogen Bubble Chambers.

The most crucial element that determines the adequacy of DAPR as an automatic scanning system is its ability to perform the track following operation. To be economically competitive with manual scanning techniques, DAPR must operate the FSD near its maximum rate, and this imposes a severe "real-time" constraint upon the program. It is our belief that if track following can be done sufficiently well within the time constraints, then all other elements of DAPR can be made to operate successfully without serious question. Conversely, if track segments of adequate quality could not be achieved within the time constraints, then no amount of sophistication in the subsequent DAPR phases would rescue the system.

Our study of the DAPR prototype results has therefore concentrated primarily upon determining the success of the track following. Examination of the prototype results shows that only an insignificant portion of any expected track has been incompletely treated. Very short tracks and those having high curvature are for the time being excluded from consideration. Closely spaced tracks should be resolved as well by the program as they can be visually on the usual scanning projector. DAPR must measure ionization at least as accurately as careful visual estimates. Within this framework we believe that the prototype abstraction program has dealt very successfully with all views that have been studied. These now amount to perhaps 50 individual views in the three chambers.

Because the prototype programs are designed to facilitate extremely detailed study of the track following procedures on one or two views at a time, no thought was initially given to system procedures which make practical runs on thirty views at one time. Furthermore, the diagnostics produced would take weeks to study thoroughly. Our experience with a representative set of data is that almost all of this is so well handled by the program as to be uninteresting, while the parts which one would like to study are almost inaccessible within the voluminous mass of detailed display.

Indeed, it is most encouraging to observe that we have outgrown the prototype techniques, because this is due to the substantial progress that has been made in the last three years. Since the production system must always retain the ability to allow detailed study of any selected tracks, it seems desirable to move now to it, rather than building volume data handling abilities into the prototype program.

Let us examine the DAPR operations from the viewpoint of data flow. Figure I illustrates this data flow by means of a block diagram.

We see that data first pass through blocks which produce the initial track abstraction. A geometric vertex search is made in each view separately, and the several views are combined during the pre-scan operations which end with the writing of a Data Abstract tape. This part of the DAPR system may be likened to the actual photographic development of the film images: the digital information is readied for use by all subsequent phases of the analysis process.

Scanning then consists of the application of physics selection criteria to the data contained on the abstract tape, so that specified events are edited into a form suited for further analysis. The primary mode of this selection is based upon an associative retrival of the data; i.e., the wanted events are described in terms of their significant properties, rather than by some label as frame number. The selected and edited events produced by this phase of DAPR are with one exception identically equivalent to those produced by conventional measuring devices such as Franckensteins, and may be processed through existing reconstruction and analysis programs.

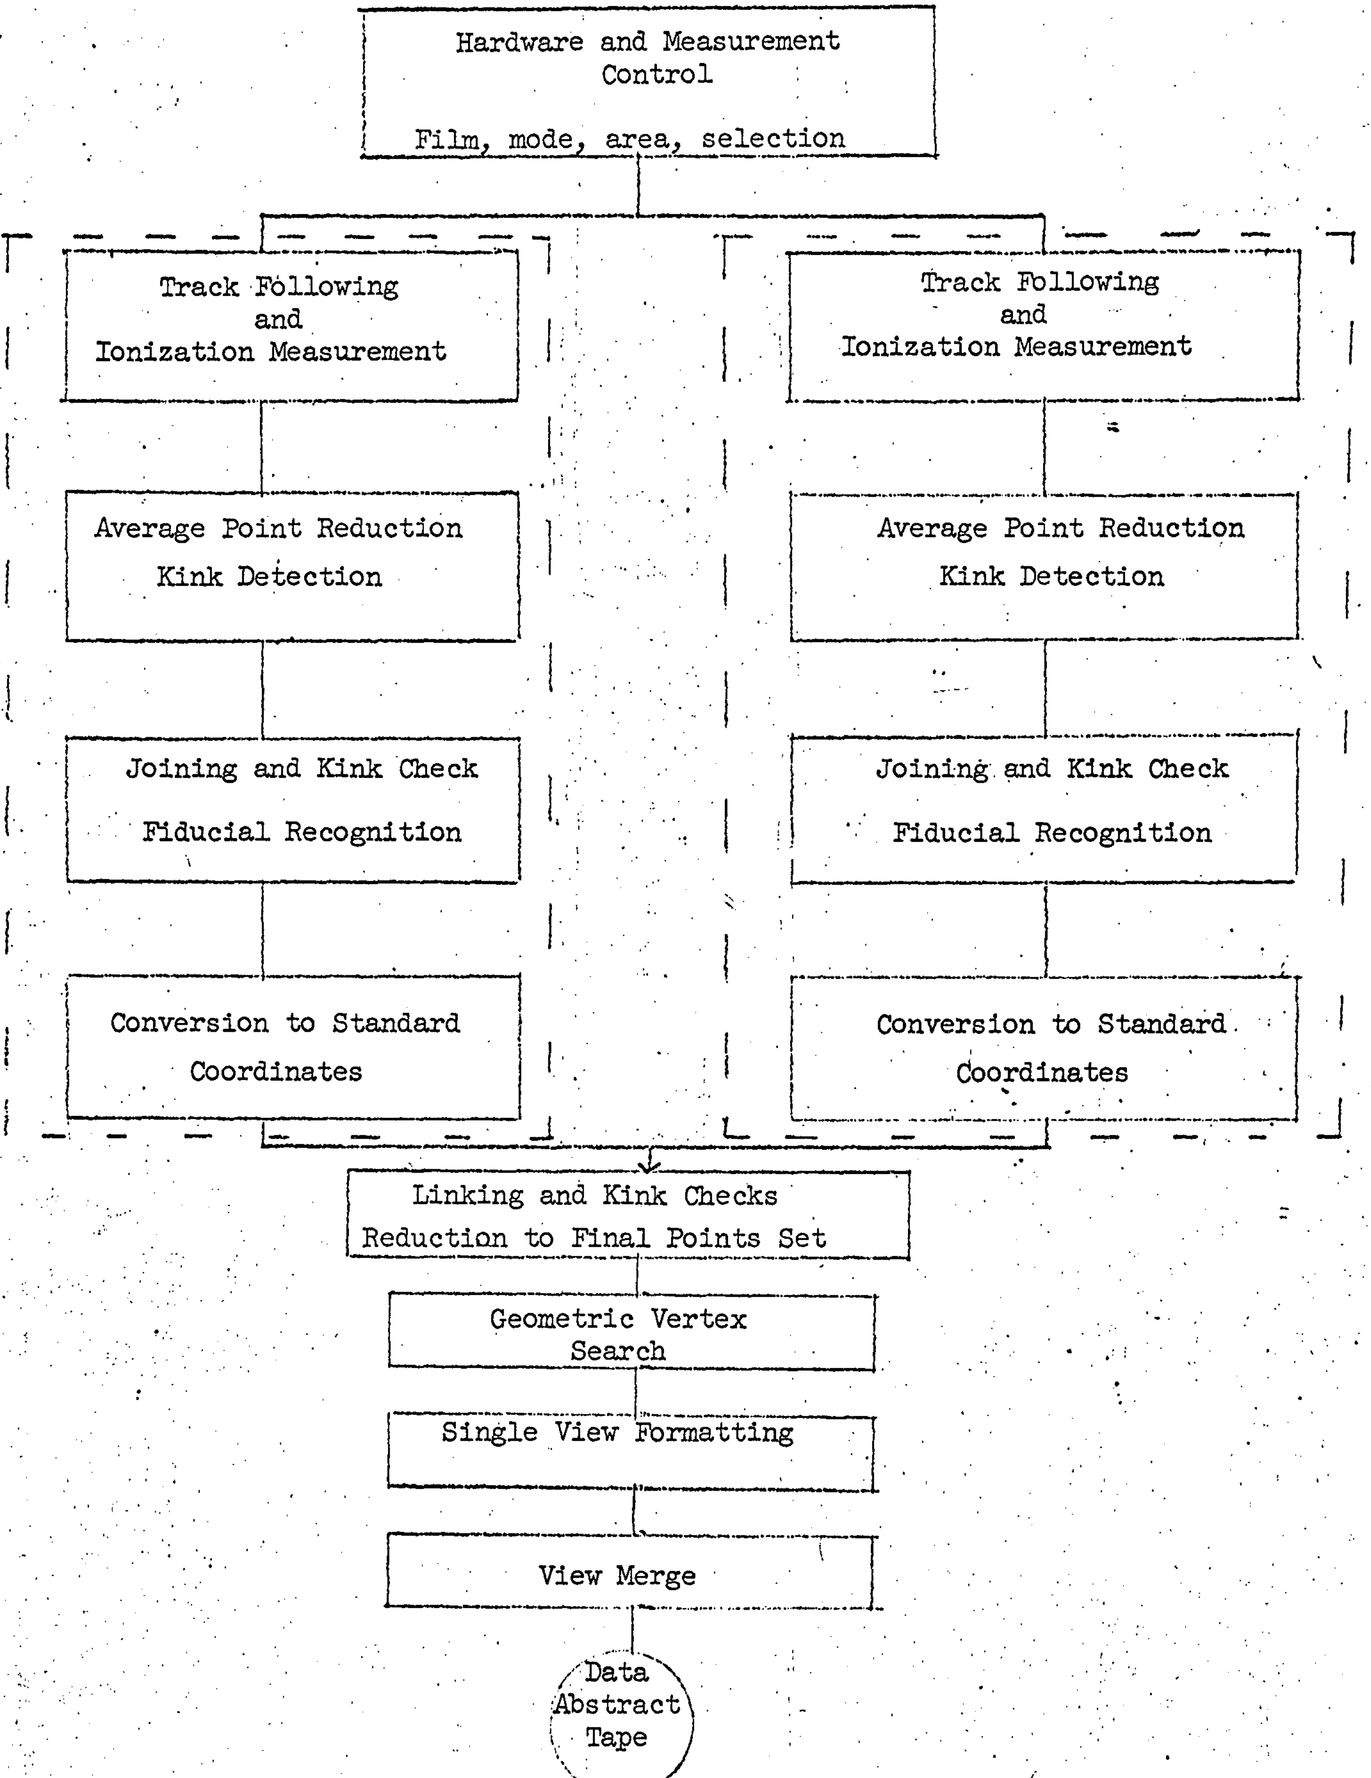

Figure I: DAPR DATA FLOW CHART

This exception is the inclusion in the DAPR data of quantative ionization measurement parameters as well as geometric parameters

We regard the prototype program as providing the kernel track following routine of the production system. This essential routine has been thoroughly studied in the prototype, and the present version will be brought forward with only a few modifications. In this crucial phase of the program the real-time demands expressed by the FSD hardware must be met. All contact with the film is made in this phase, so that all further processing depends upon its successful completion.

A detailed evaluation of the several hundred tracks in the views so far analyzed indicate that the current procedure is effective in following tracks of all configurations. Some departures from ideal operation remain and will have further improvement before reaching the final system. These problems lie in three categories which are:

1. The inclusion into average points of digitizings lying within the roads but not on the tracks.
2. The unflagged depletion of digitizings in non-beam tracks within areas being crossed by beam tracks, and the subsequent systematic bias decreasing the measured ionization.
3. The failure of the prototype program to recognize tracks having fewer than 12 digitizings, and thus the categorical rejection of tracks shorter than about 1.5 centimeters in the chamber.

Several possibilities exist as solutions to each of these problems; and a final choice will be made on the basis of early experience with the production programs.

The prototype programs have also yielded a very satisfactory kernel for the track joining and linking phases of the production system, and will be used with only minor changes. These routines do not have the real-time constraints which are applicable to the track following routines, and therefore are much less demanding of programming sophistication. Examination of the data so far obtained indicates a very good correlation between actual tracks in the picture and the final segments resulting from application of the joining and linking routines.

The geometric vertex search routine, which merely attempts to find in individual views the points representing common intersections of track segments, has recently been appended to the prototype program. An enlarged and more capable version incorporating the same general procedures will be made a part of the production system.

Because we have not regarded the problem of event selection as being different from one which is already being satisfactorily handled in bubble chamber analysis programs, we have not thought it desirable to spend programming effort on prototype programs of this type. It is evident that if track abstraction data of sufficient quality are at hand, straight-forward application of programming techniques can readily be made to impose physics selection criteria upon the data. Track correlation between several views is already being performed by computer in Franckenstein and Spiral Reader data. Event analysis programs for use with Franckenstein and HAZE-FSD systems are directly useful to the DAPR system without modification.

For DAPR to be truly a contribution to bubble chamber data analysis technique, it must actually perform scanning and measurement of bubble chamber film with at least equal quality and lower cost than conventional methods. We have not doubted that automatic procedures could meet or even improve upon the quality standards set by conventional methods, if unlimited amounts of computer time are made available. Rather, we have wondered whether both quality and cost standards could be met with existing hardware. The prototype programs have demonstrated that they can meet these demands, and can operate the Berkeley FSD at approximately its maximum rate.

Table I summarizes the times required by the FSD to carry out various parts of the measurement with film from the different chambers.

Table I

Single FSD Measurement Times for Film from Various Chambers*

Chamber	25" HBC	72" HBC one orthogonal scan	72" HBC two orthogonal scans	80" HBC two orthogonal scans
Picture area, mm	45 x 40	125 x 40	125 x 40	150 x 50
<u>Times</u>, sec.				
Move film } overlapped	1.9	2.7	2.7	2.9
Retrace stage	0.7	1.0	0.5	0.7
Sweep normal	1.7	4.4	4.4	5.2
Move stage orthog.	0.7	0.7	0.7	0.8
Sweep orthog. No. 1	1.5	1.5	1.5	1.9
Move stage orthog.	-	-	0.7	0.7
Sweep orthog. No. 2	-	-	1.5	1.9
Output	0.2	0.2	0.2	0.2
Total	6.0	9.5	11.7	13.6
Active Computing	3.4	5.9	7.4	9.0
Fraction Active	0.57	0.62	0.63	0.66

---

*Assumes consecutive views scanned.

Our present FSD model is composed of one film-measuring unit and a set of digital electronics. The consequence of this arrangement is that we have available only a limited degree of simultaneous operation. It is possible to time-share the film move operation with the major stage retrace, but otherwise all operations must proceed sequentially. Addition of a second film-measuring unit to the existing system makes possible a very extensive amount of time-sharing with relatively small increase in hardware cost. We call such an arrangement a tandem unit (TFSD), and we hope to build such a machine at Berkeley before long.

It is desirable to study in some detail the maximum measurement rates that can be achieved with a TFSD. For this it is useful to construct a chart in which the various operations necessary for time-shared measurement of two bubble chamber films are shown in their correct time sequence. Figure II is such a chart, in which we have selected the 72" HBC film with doubled orthogonal scan as being intermediate and thus representative of the film types which we can presently use.

In Figure II, we apply the times taken from Table I to the various operations required to obtain a measurement. Note that we here assume that the computer and its programs are able to keep pace with data generated by the TFSD, which assumption defines DAPR as being a real-time program. It is evident that the TFSD maintains a rather efficient duty cycle with respect to the computer, since data are sent to the computer during 15.2 of each 16 seconds. In these 16 seconds, two complete views are measured, yielding an implied rate for 72" HBC film of 150 triads per hour even with the double orthogonal scan.

The prototype programs have demonstrated that the track following mode of DAPR can indeed keep pace with the TFSD. We recall that, for normal mode scans, the FSD digitizes the Y-motion of the stage at the beginning of the sweep, the X-motion of the stage at the beginning of each scan line, and the W-motion of the spot at each intersection of the scan line with a track. Orthogonal mode scans are digitized in a similar manner. During the time of data input to the computer, one of the FSD units is sweeping 480 scan lines per minute, and our experience shows that on the average 15 - 20 "W" digitizings per scan

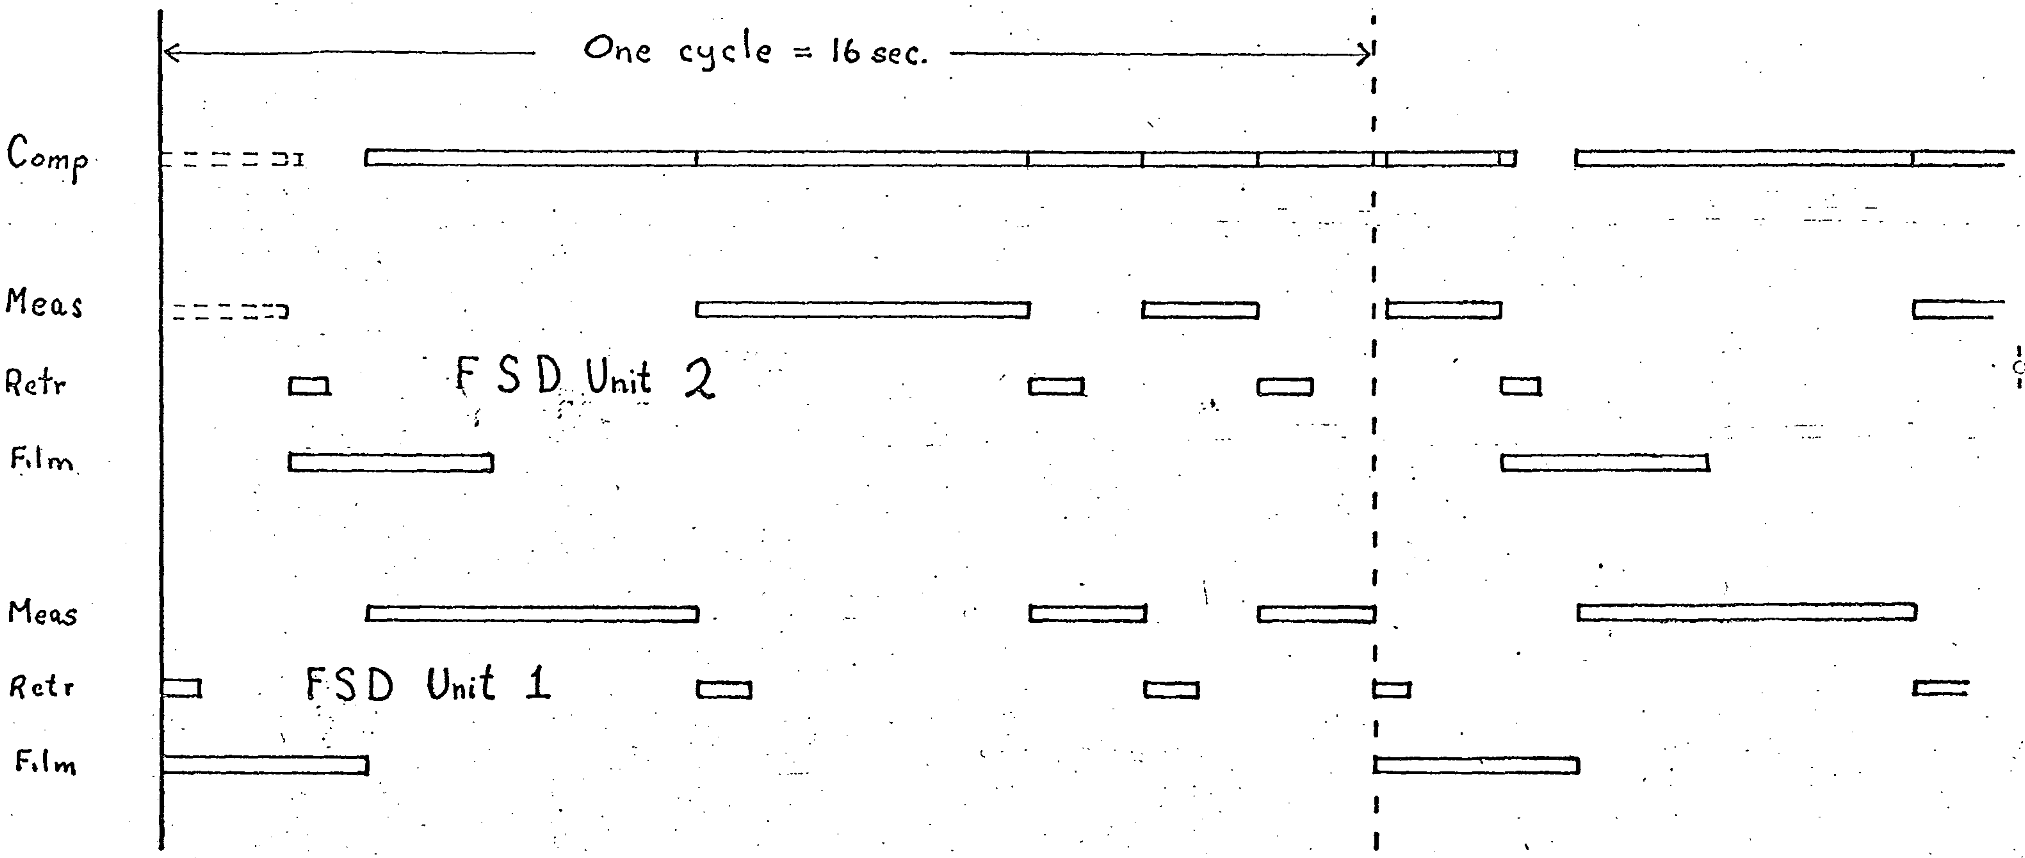

Figure II: Time Shared Operations with TFSD

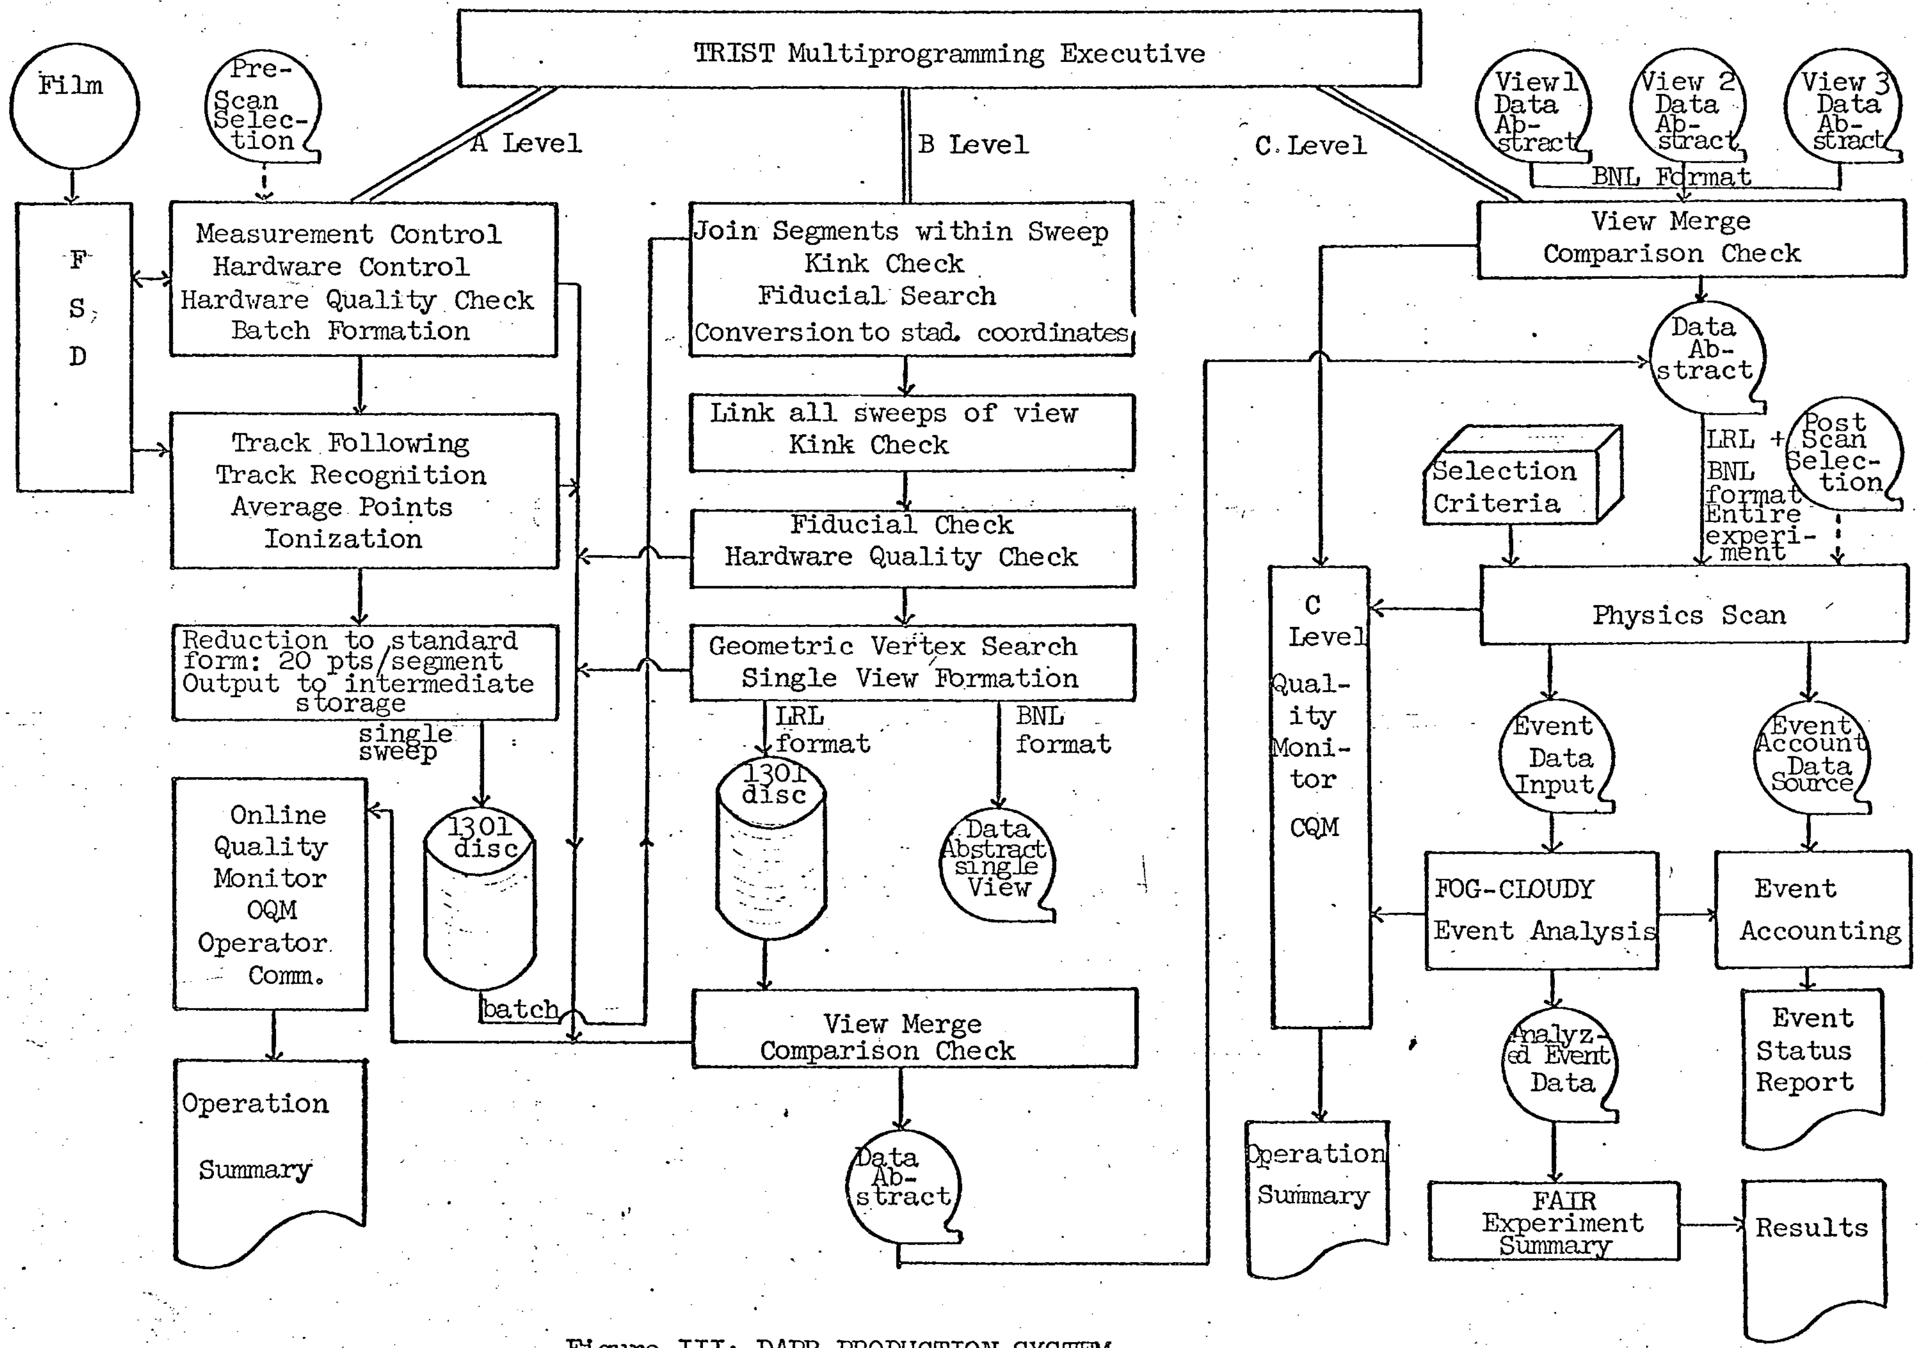

Figure III: DAPR PRODUCTION SYSTEM

line are to be had with almost any film. Thus about 8,000 "W" digitizings per second are delivered to the computer. The prototype DAPR track following routine handles on the average more than 14,000 "W" digitizings per second. Although this 50% safety factor may seem large, it must be remembered that fluctuations in the data rates require instantaneous capacity greater than the average.

It may be argued that a newer and larger computer would more easily meet this real-time demand but the criterion of least cost per event while maintaining quality must be considered. Use of a computer which greatly exceeds the need causes some difficulty in adequately using the remainder of its capacity, since most operating monitors do not have a sufficiently well developed real-time response capability. Further, if one must do substantially more computation per event, the cost per event rises proportionately, and at some point defeats the improved cost per computation ratio of the newer computers. We do not know of another existing computer that would yield a lower event cost to us than the laboratory owned IBM 7094 II. This should not be taken to mean that we do not plan moving to a newer computer at some time in the future. Rather, we feel that this production DAPR system will be a useful tool of physics, and will give us experience leading to the formulation of the next generation system.

The production DAPR system is therefore being implemented on the LRL IBM 7094 II computer. This computer has a total of five data channels, including three which connect sixteen tapes, a printer, a punch, a card reader and a CRT. The other two channels are assigned separately to the Direct Data Connection, and to a disc file unit. Memory consists of two banks of 32,768 words each, overlapped in the usual manner for the IBM 7094 II. The standard multiprogramming package, including memory protect, is installed on this computer.

TRIST, a multiprogramming executive has been especially written to allow the real-time FSD programs to share the computer with other data analysis programs, so that full use can be made of the entire computing capacity. This executive has been in routine use with the HAZE system since February 1965.[5] It has provision for three levels of program priority, and as many as fourteen programs may be in shared operation at one time.

The highest priority is assigned to the FSD real-time program, which in the case of DAPR contains the control and track following routines. This program is continuously resident in one of the two memory banks, together with the TRIST executive itself. Data from the FSD arrive under control of a dedicated I/O channel. When this channel exhausts its command list, a trap to the central processing unit must be serviced in time to prevent loss of data. TRIST meets this request within less than 500 microseconds, and allows the priority program to have complete control of the computer for whatever period of time it requires.

The second priority level is assigned to a stack of as many as twelve programs which batch process a small number of events just after their measurement. In the case of TFSD operation, two logical stacks of as many as six programs are used. In the case of DAPR, these programs perform the joining, linking, and vertex search operations. A batch is formed under control of the measurement program ("A" priority level), and the executive is notified of the exact sequence of secondary priority programs ("B" priority level) which is required. These are called into core one by one, and do their processing operations during the central processor time which is not used by the "A" level program. Each program initializes itself when it receives the first batch of events, and remains in a state of interrupted operation between batches until the last, when it goes through the usual close-out operations. Data pass from one program to the next through the disc file, which is used as intermediate storage.

The third priority level is assigned to a background data processing program which operates on data measured at some other time, and which uses the central processor during the times when "A" and "B" level programs do not need it. Since to this "C" level program the TRIST executive looks just like the older FCF executive, any of our bubble chamber data processing operations can be run in this way, including computations performed on Franckenstein measured events.

The TRIST executive calls each of the "B" and "C" level programs into the second core bank from the disc. All words of core, and all registers are restored to be the values which they had upon exit from that program in its last visit to core, and processing is resumed. Control is switched between the "A" and the "B" or "C" level programs until a different program is required in the second core unit. When this happens, the executive allows input/output operations to cease, and stores the entire contents of the core memory and the machine registers on the disc to await the next call of this program. "B" level programs are not retired to the disc until they have completed a batch of events, but the "C" level program may be interrupted and retired with control at any arbitrary location.

The DAPR production programs may readily be organized into a form that makes maximum effectiveness of this multiprogrammed computer environment. The FSD control and track following routines become the "A" level program, the joining, linking, vertex finding, and editing routines become the "B" level programs, and the scanning and analysis operations are done by "C" level programs. The scheduling goal is to have the "A" level program operate the FSD or TFSD hardware as nearly continuously as possible, to have the "B" level programs process batches of events whenever they accumulate, and to use whatever computer capacity remains for "C" level production work.

Figure III illustrates the DAPR production system organization. The configuration is shown as it will be when DAPR first begins physics production about July 1966. The TFSD will not be operational at that time, so the figure shows a single FSD. Provision is made for film of both LRL and BNL formats. It is seen here that the LRL format is substantially better for this use, since it allows the comparison of views to be made in the on-line phase of the process, and thus allows a greater significance to be had from the quality control information.

The "A" and "B" level programs are concerned with the on-line measurement, and their objective is to produce the Data Abstract tape. This tape is desirably a complete digital abstract of the film, containing in digital form all useful information and none of the noise found in the actual picture. The programs which produce it conscientiously refrain from making use of any information except that which can be measured in the film, so that the tape represents an unbiased measurement of the film data.

For two reasons it is desirable to put as much as possible of the on-line measurement programming into the "B" level. The memory available for the "A" level program is only 24,576 words, while several core-loads of nearly 32,768 words each may be used by the "B" level programs. The "A" level program must maintain exact pace with the FSD hardware, while the "B" level programs can average the statistical fluctuations in the data of several pictures.

The "A" level program must accomplish the primary data reduction process. Typically 35,000 "W" digitizings are input to the computer during a normal scan of 72" HBC film. These describe the ten fiducials, perhaps fifteen beam tracks, and usually a total of less than fifteen other tracks, together with optical noise from many sources. Desirably these would be reduced to perhaps 300 words of high significance that would contain all useful information in the picture. Thus, it is required to apply a reduction factor of 100 to the data. It would be much too wasteful of computer time to attempt to save all data, since this would require the writing and then reading of massive amounts of data. It is true that this can be done on modern computers simultaneously with other operations, but the usual difficulty with this approach to data processing operations is that a much larger memory is required for a staging area than current computers possess. It is better then to make a substantial reduction as soon as possible.

A lower limit on the number of points which need be saved at this phase is imposed by the need to search each segment for the possible presence of a "kink". The track following process is such that two physical tracks meeting at a small angle may not be distinguised in the digital output, but may yield one segment containing track elements from both sides of such a scatter or kink. Our experience in seeking kinks in data measured by Franckensteins is that almost a continuum of points along the track is needed, but we believe that the greater accuracy and more uniform distribution of points in the track segments yielded by DAPR makes kink detection practicable with as few as twenty points.

The "A" level program outputs to the disc intermediate storage a set of twenty points spaced uniformly along each followed track segment. Actual tracks may be followed in their entirety, or may be represented by two or more segments due to being lost and then re-initialized in the track following process. Each point is the average of four consecutive digitizings on the track.

Ionization measurements are obtained for each segment. This measurement is the ratio of digitizings made on the track to the total number of intersections of the flying spot with the track. It is measured only over the portion of the track segment where no interference from crossing tracks is recognized.

The "A" level program schedules the sequence of hardware operations and defines the batches of data for use by the "B" level programs. The scheduling algorithm is formulated in accordance with a diagram similar to Figure II, and has as its goal achieving the maximum density of central processor usage while still meeting the real-time constraints. In order to keep the frequency of interchange between programs of the "B" and "C" levels at a reasonable level, the batch size is chosen so that one batch will be produced each five minutes. The control program organizes the required number of consecutive views into a batch, and when all have been stored on the disc, a flag is set to cause the TRIST executive to initiate a "B" priority cycle. A 72" HBC batch consists of forty-five views.

The batch is processed by a cycle of "B" level programs to yield the data abstract tape. The production system is quite similar to the prototype described in the previous paper, except that more attention must be given to the presence of kinks. The first operation performed is therefore a search for kinks within the track segments produced by the track following program, and after each of the joining and linking operations a further check is made for kinks especially at the point of joining.

A major difference between the prototype and production programs is the inclusion in the latter of an extensive checking routine, designated the Online Quality Monitor (OQM) in Figure III. Our experience with the HAZE-FSD system has shown that one of the most difficult tasks in an automatic measurement system is the immediate discovery and correction of hardware and operator errors. HAZE depends upon manually generated roads to define events to be measured, and the agreement between the roads and data measured on the FSD provide a sort of redundnacy check for accuracy. DAPR has no such external check, so that entire dependence must be placed upon checks internal to the track abstraction phase. Not only must hardware malfunctions be prevented from degrading the data, but a wide variety of problems encountered in the film must be recognized and properly dealt with.

The OQM monitors a series of hardware signals which are set by the occurance of errors in several fundamental operations of the FSD. These include checks on each of the three coordinate measurements, so that a slip of zero point or change of scale factor would be immediately detected and made known to the computer. Film move commands are checked, and those which cannot be executed because of film difficulties are reported back to the computer. Should the real-time program fall behind the rate of measurement, and thus lose data, the FSD hardware signals this to the computer. All of these hardware checks are tested for each scan of the film, and the occurance of any error indication causes the control program to attempt corrective action.

A major cause of rejected events in HAZE has been the difficulty of positioning a required film view for measurement. Several factors contribute to this: incorrect data box markings on the film, unreadable markings, splices and film tears, as well as failure of the FSD to reach an otherwise valid film view. In many cases the FSD hardware is aware that it is unable to position the requested view, and sends the proper signal to the OQM. However a significant portion of the positioning errors are not recognized by the hardware, and require discovery by the computer program from information transmitted by the measurement.

Internal checks for proper film positioning can be made on single views only on the basis of relative location of fiducial marks. This is possible because most chambers have at least some fiducials out of the primary fiducial plane, and therefore allow determination of the camera location. Film in the LRL format, which has all three views on a single film, offers a substantial advantage in this regard, since most unrecognized positioning errors result in transmitting measurements of the adjacent view. The adjacent view is made by a different camera, and therefore a comparison of fiducial locations with those expected immediately signals an error.

Comparison between the several views of one picture offers still more checks that each view is of the common picture. We find that the DAPR beam track count is quite accurate, so that comparison of number of beam tracks in the candidate views is useful. Such a check is not absolute, because beam counts are often perturbed slightly by tracks which appear coalesced during their entire passage through the chamber when seen from one camera angle. However, in many beams, the fluctuation in the number of beam particles per picture is very much greater than inaccuracies in beam count.

The obvious check which must be satisfied if all views describe the same picture is made by comparing geometric vertices. This comparison must recognize that many of the geometric vertices will be accidentals of one view only, and that some of those which are common to all views will appear differently in each view. However, if an actual physical event occurs in a picture, it gives a very powerful evidence for identification of the several views. This process is really very much like the one which a person scanning a film goes through; one immediately recognizes that a stereo set is not being viewed by seeing a lack of correspondence between the several views.

Since it is possible for film system hardware to lose contact with the film during a run of reasonable duration, it is extremely important to have these checks occurring continuously throughout the run. This is possible only with LRL format film, and is the reason why there are two blocks labeled "View Merge" in Figure III. Although the process is almost the same in the "B" level program for LRL format as it is in the "C" level program for BNL format film, the use to which the comparison checks can be put is entirely different. In the case of the "B" level run, reports are made to the OQM, and valid operator action can be taken soon after contact with the film is lost. If comparison is deferred until "C" level, it would be possible for an on-line run to go to completion without the operators being aware that film contact had been lost, and, of course, a major element of confusion would have been introduced into the system.

We assume that most of the DAPR system use will be in measurement of events having high density on the film. For film having an event in every one or two frames, the system will be economical and will not require manual prescanning or selection of any kind.

It may be desired to preselect certain frames containing desired events in the case that the experiment being performed has sparsely distributed events. Such preselection of frames containing relatively rarer events would presumably not bias the abstraction of more common events, and may well be the manner in which experiments of medium and high frequency events are combined. An optional form of data entry is shown in Figure III, and appears as input to the measurement control section of the "A" level program. This tape may be generated on the SPVB's used by the HAZE system by measuring one fiducial and the production vertex in one view for each event wanted, and may include the event identification code. The input tape may also be generated by keypunching only a list of frame numbers. In any case, the control program makes use of this input if proper options are elected before the run begins, and attempts to produce a data abstract tape only for those frames mentioned in the special input. In the case that this option is not selected, data contained on the special input selection tape are carried through to the data abstract tape in parallel with the track measurements, so that the later scanning programs can make use of this additional source of information. Provision is also made for using a selection tape as input to the physics scan program, so that post (automatic) scanning selection can be performed. One very useful purpose of this feature is to allow easy comparison of manual and automatic scanning results, since a simple listing by the event accounting routines can show comparisons between these two scanning results on an event by event basis.

Once a data abstract tape has been generated for a sufficient quantity of film, the actual scanning process can begin. Data are stored in very compact form on the data abstract tape, and therefore a substantial volume can be scanned during a short run on the computer. We expect that one abstract tape will contain all track data from perhaps 3500 triads, and that it can be searched in about ten to fifteen minutes of 7094 time.

Selection criteria are read by the scanning program, and a simple compilation process yeilds a special program which will perform the required scanning tests. The selection criteria are no more than somewhat stylized scanning instructions, and should be no more complicated to write. As is usual for manual scanning, a simple set of instructions will be written, and then as a little scanning takes place, corrections and further specifications will be added. We expect that the same process will occur in the automatic mode, except that after each change it will be feasible to go back to the beginning of the experiment, and thus to have a homogeneous experiment when the final change has been made.

It is expected that a fairly heavy dependence will be made upon the ionization measurements, just as manual scanning gives high weight to this information. Relative ionization is useful not only in defining mass possibilities, but also is a useful tool for correlating tracks in the several views.

Output from the scanning process will be a tape which contains the set of points for each track in each view, the ionization measurement, the track numbers assigned by the scanning procedure, and all information usually yielded by any other film measurement process. The difference is that this part of the DAPR process can run nearly at tape reading speed through the film abstract. Conventional programs which reconstruct and analyse Franckenstein measurements can work directly with this output tape.

We expect that certain features of the Three View FOG program will be of great help in determing the quality of the DAPR results[6]. Hough[7] has pointed out the importance of the redundancy contained in a third view as a check on the quality of measurement, and our FOG program makes quantative use of the residuals from the best orbit projected into each view to determine the quality of the measurement. It is expected that many of the further improvements in the DAPR techniques will come from study of difficulties turned up by such techniques.

We look forward with some confidence toward the day when DAPR becomes a major tool of bubble chamber physics. The prototype programs have shown that film can be abstracted suffieciently well and economically to provide a good basis for digital scanning. Experience with the measurment of 250,000 events in the HAZE-FSD system has demonstrated that the hardware and multiprogramming executive perform reliably, and has pointed to some pitfalls to be avoided in this new system. We believe that experience with HAZE has stimulated the experimenters to strive more diligently for film amenable to automatic measurement. If all goes well, we expect to come to next year's FSD conference and tell results of an actual physics experiment.

References

1. Charles Dickens, Mary Downton, Joan Pinchak, and Howard S. White. "Results from Prototype DAPR Programs", Lawrence Radiation Laboratory, Berkeley, UCID 2653. Paper delivered at Columbia FSD Conference, 1965.

2. P. V. C. Hough and B. W. Powell, Proc. Int. Conf. Instr. H. E. Phys., Berkeley (Interscience Publishers, N. Y., 1961).

3. H. S. White, et. al. Preliminary Operating Experience with Hough-Powell Device Programs. Nuclear Instruments and Methods, North-Holland Publishing Co., 1963, pp. 393 - 400.

4. H. S. White. "Status and Future Plans for LRL Flying Spot Digitizer." Lawrence Radiation Laboratory, Berkeley. A paper delivered at the High Energy Instrumentation Conference, Purdue University, May, 1965; IEEE Transactions on Nuclear Science, NS-12, 4 (1965).

5. C. Osborne, R. Larson, T. Oliver, J. Stekler and H. S. White. "TRIST, A Tri-level Multiprogramming Executive Involving Real-time Calculations and an Asynchronous On-line Data Source." Lawrence Radiation Laboratory, Berkeley, UCRL, 16199.

6. Dennis Hall, et al. "The Multi-view Fog Program and its Application to Quality Control of FSD Data." Lawrence Radiation Laboratory, Berkeley, UCID 2652. A paper delivered at Columbia FSD Conference, October 1965.

7. P. V. C. Hough and B. W. Powell. "On Getting Results from an HPD System." Paper delivered at Purdue Conference on Instrumentation for High Energy Physics, May, 1965; IEEE Transactions on Nuclear Science NS-12, 4 (1965).