BNL 725 (T-259)

BROOKHAVEN LECTURE SERIES

A Computer Learns To See

Paul V.C. Hough

Number 14

February 14, 1962

BROOKHAVEN NATIONAL LABORATORY

Associated Universities, Inc.
under contract with the
United States Atomic Energy Commission

FOREWORD

The Brookhaven Lectures, held by and for the Brookhaven staff, are meant to provide an intellectual meeting ground for all scientists of the Laboratory. In this role they serve a double purpose: they are to acquaint the listeners with new developments and ideas not only in their own field, but also in other important fields of science, and to give them a heightened awareness of the aims and potentialities of Brookhaven National Laboratory.

Before describing some recent research or the novel design and possible uses of a machine or apparatus, the lecturers attempt to familiarize the audience with the background of the topic to be treated and to define unfamiliar terms as far as possible.

Of course we are fully conscious of the numerous hurdles and pitfalls which necessarily beset such a venture. In particular, the difference in outlook and method between physical and biological sciences presents formidable difficulties. However, if we wish to be aware of progress in other fields of science, we have to consider each obstacle as a challenge which can be met.

The lectures are found to yield some incidental rewards which heighten their spell: In order to organize his talk the lecturer has to look at his work with a new, wider perspective, which provides a satisfying contrast to the often very specialized point of view from which he usually approaches his theoretical or experimental research. Conversely, during the discussion period after his talk, he may derive valuable stimulation from searching questions or technical advice received from listeners with different scientific backgrounds. The audience, on the other hand, has an opportunity to see a colleague who may have long been a friend or acquaintance in a new and interesting light.

The lectures are being organized by a committee which consists of representatives of all departments of the Laboratory. A list of the lectures that have been given and of those which are now scheduled appears on the back of this report.

Gertrude Scharff-Goldhaber

The drawing on the cover is taken from a 5th Century B.C. relief on the Acropolis in Athens, the "Dreaming Athena," by an unknown sculptor.

BNL 725 (T-259)
(Mathematics and Computers – TID 4500, 17th Ed.)

A Computer Learns To See

Paul V.C. Hough

UNITY OF SCIENCE * BROOKHAVEN LECTURE SERIES *

BROOKHAVEN LECTURE SERIES

Number 14, February 14, 1962

PRINTED IN USA
PRICE 50 CENTS

Available from the
Office of Technical Services
Department of Commerce
Washington 25, D.C.

June 1962 1225 copies

INTRODUCTION

Only 12 years ago high energy physicists went up high mountains with their cloud chambers, emulsions, or Geiger counters and contentedly caught what results they could. The beams of the particles they were investigating were not at all well collimated and contained particles that only rarely produced anything of interest. For instance, with cloud chambers, in which liquid drops condensed along the paths of charged particles were photographed, a few interesting events were obtained per day. A good projector or a microscope, some knowledge of geometry, and a slide rule were quite sufficient to keep up with the analysis of the photographs.

When the big accelerators, the Cosmotron at Brookhaven and the Bevatron at Berkeley, went into operation, measuring and computing of events started to fall behind and some efforts were started towards automation, but cloud chambers were so inadequate to match the accelerators that the need was not urgent. Then Donald A. Glaser at the University of Michigan invented the bubble chamber, in which photographs are taken of bubbles produced in a superheated liquid along the particle paths; and the production rate of events of interest went up by two orders of magnitude, since, because of their high density, liquids are very effective targets. As bubble chambers grew in size, measuring and computing fell far behind, and this part of high energy physics has become rather unsatisfactory in spite of much progress in automation.

Going back to the early 1940's, there was at Swarthmore College near Philadelphia a very successful student in physics whose name was Paul V.C. Hough. Immediately after graduation he went to Los Alamos, from where he came back to Cornell University to receive his Ph.D. degree in 1950. Next he went to the University of Michigan, where he became an Associate Professor in 1956. As mentioned, Donald Glaser, the inventor of the bubble chamber, was also at Michigan during that time, and there must have been some inspiring discussions as to what to do with bubble chambers when they grew up.

A few examples of Dr. Hough's publications are as follows: "The angular distribution of pair-produced electrons and bremsstrahlung," "Radial oscillations in the cyclotron – experimental," "(d,t) reactions and the triton wave function," "States of low excitation in O^{18}," "Cyclotron instrumentation for nuclear spectroscopy at medium resolution in energy."

Starting in 1949, some others are "Automatic grain counting of tracks in nuclear emulsions," "Nuclear emulsion scanner," "Scanner recognizes atomic particle tracks," and, finally, in 1960, "A method for faster analysis of bubble chamber photographs," a very timely development of great interest in view of the aforementioned difficulties.

We were most fortunate that Dr. Hough came to Brookhaven early in 1961, bringing along all his energy and enthusiasm. He will now tell you about his work which, I believe, will give high energy physics and perhaps other fields a tool that will make it possible to do more and better work leading to new interesting results.

RALPH P. SHUTT

A Computer Learns To See

1. COMPUTERS AND HUMAN VALUES

Figures 1 and 2 show the IBM 7090 computer at Brookhaven – one of the collections of hardware which are causing surprising social changes in the world today, especially in the United States. Most of us have seen pictures of acres of chemical processing plant, barren of people. The long rows of desks in accounting and payroll departments are disappearing. The rivers of parts flowing together to form that rather stagnant sea known as the automobile are now often controlled by a central computer. Electric power generation and distribution are beginning to come under computer control. Serious thought is being given to computer control of the operation of the 80-in. bubble chamber at Brookhaven. The manufacture of small electronic parts and even steel making are becoming self-regulating systems. This "second industrial revolution" is having a remarkable effect on our civilization. However, there is a second perhaps even more profound computer development. Computers are having an influence on human values, and this area I want to explore a bit here.

From time to time someone wants to try to get some help from computers in areas of human activity which are not very routine. For example, the scanning and measuring of bubble chamber pictures is quite a complex human operation, but there is so much of it to be done that we are trying hard to turn large blocks of the work over to our 7090. When we try to get a computer to solve a nonroutine human problem, two things happen.

First we learn much more about the problem. In fairly mysterious problems, like the recognition of shapes or patterns, a sequence of elemental computer operations which will *recognize* amounts almost to an operational *definition* of the process of recognition. We begin to understand what objective criteria have sufficient power to achieve discrimination among patterns. We discover how much information (and what kind) must be transmitted and evaluated. The very limitations of computer intelligence force us to break each overall human function into elementary subfunctions. The consideration of these subfunctions together amounts then to very searching analysis of the process.

The second thing that happens is that we gain new insight into human beings. At first this new insight always seems faintly deprecatory. That the computer can perform some task we have prided ourselves in is unsettling. However, if we just accept such classical human limitations as low precision, low speed, and fluctuating motivation, we can (with the aid of the also classical human device of a sense of humor) regain our unperturbed state and consider objectively the human performance. Our curiosity is aroused as to whether the elementary computer operations which solve the problem are duplicated by the human. If not, what are the human subfunctions? How is information transmitted and stored in the human nervous system? But, beyond these interesting detailed questions, we are led to inquire: "What are the peculiarly human capabilities? What are the human functions which are in essence not transferrable to a computer?" This is not the kind of question for which a pat answer is possible, or even valuable. But answers will begin to be clear over the next years and decades and will, I believe, have a gradual, powerful influence on human values. This influence will be similar to and perhaps comparable with that from the detailed understanding of genetic reproduction, and genetic control of protein building, which now seems with us.

2. WHAT IS A COMPUTER?

Let us return now to the computer itself. What is a computer? A computer is essentially a device for storage and manipulation of numbers. Let us take a moment to see how it does that.

Very much of a computer is simply a collection of elementary physical systems, each of which can exist in two states. Call them "binary systems." The two most common types nowadays are shown in Figure 3.

The left half of the diagram shows two transistors, each of which can carry a current I_0 or no current at all. The crisscross lines between the

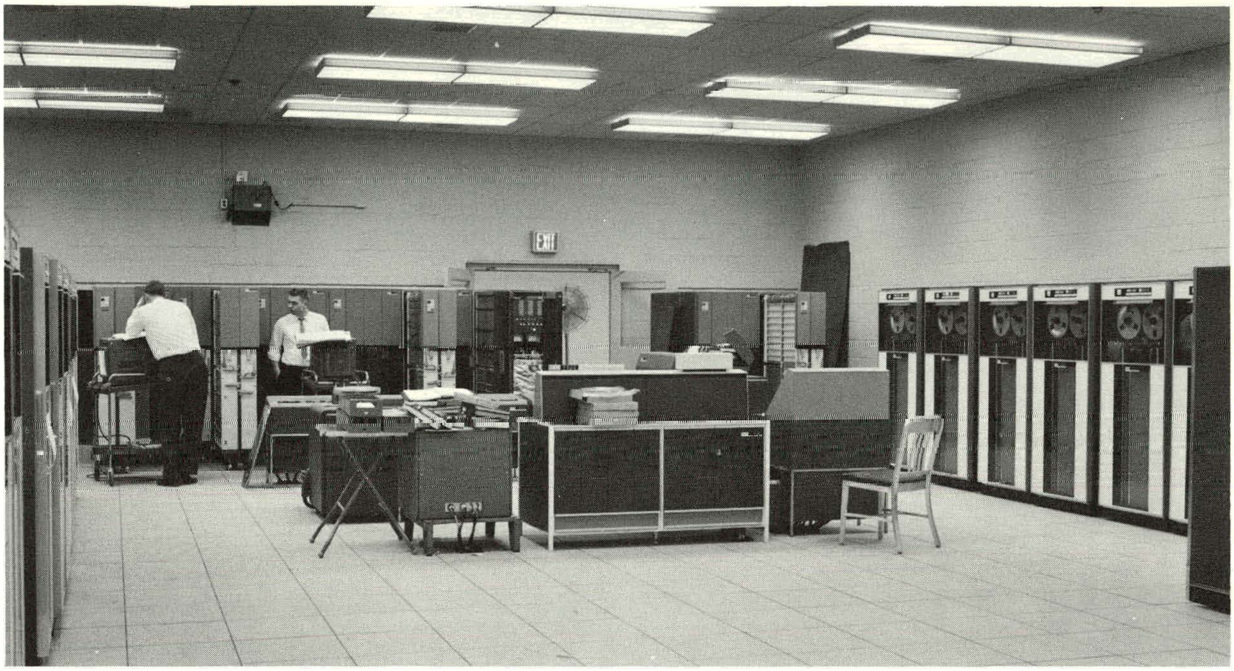

Figure 1. A general view of the IBM 7090 computer installation at Brookhaven.

Figure 2. A view of the 7090 showing the control console.

transistors indicate the interconnection of the transistors, which is such that when the left transistor is conducting the right is inhibited and also vice versa. Evidently the effect of such an interconnection is to "lock in" the state "left transistor conducting, right transistor not" and also the reverse state. A transition between the two states requires a forcing current injected from outside the system.

The right half shows a small doughnut of ferromagnetic material carrying magnetic induction B in either the counterclockwise or clockwise sense. Here the "lock-in" to one or the other of these states arises from the interaction among the electrons in the solid, the cause of the basic phenomenon of ferromagnetism. Again a transition between these two stable states requires forcing from outside the system, here in the form of a magnetic field generated by a current-carrying wire which threads the doughnut.

What good are these binary systems? Consider a couple of numbers:

$$13 = 1\times8+1\times4+0\times2+1\times1$$
$$= 1101$$
$$5 = \qquad 1\times4+0\times2+1\times1$$
$$= \ 101$$

Now use the state of a binary system to represent a 1 or a 0. Then

for example.

This is how a computer stores numbers. [Just by the way, it is possible to make a physical system with 10 states, and so represent decimal numbers, but such systems are far more complicated and hence both more expensive and less reliable. As far as I know, all modern computers are binary inside.]

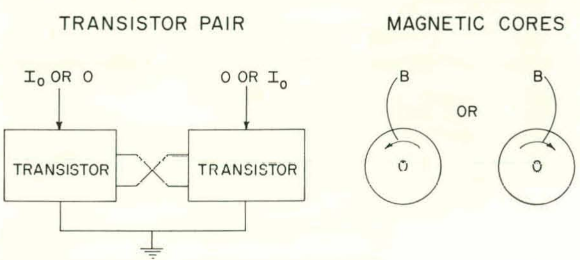

Figure 3. The two most common binary systems in present day computers.

How can the computer manipulate numbers? A perfectly typical example of manipulation is ordinary addition:

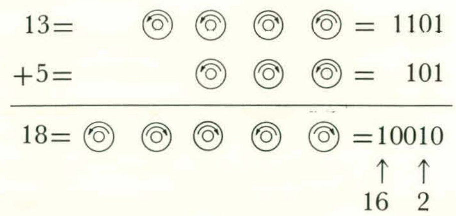

The rules for carrying out the addition are easy to see from the meaning of the binary digits. Consider the far right column. The meaning of the two "1's" is "number of units." Two units add to one two with no units left over so the result of the addition is a "1" in the next column (called a "carry") and a "0" in the far right column. If there had been only one "1" in this column of the addends it would be preserved in the sum, without causing a carry; and of course the sum of two "0's" is a "0" with no carry. To carry out the addition electronically inside the computer requires only the generation of voltages representing "1's" and "0's" at terminals representing the columns of the sum, by using circuits which faithfully carry out the rules we have just found. It turns out that these circuits are made up easily by computer engineers.

3. HOW DOES THE COMPUTER DEAL WITH THE WRITTEN LANGUAGE?

For better or worse, there is nothing profound here. A one-to-one numerical code is set up for A - Z and for a set of useful signs like *, /, @, space, etc. Of course the computer has to know whether the numbers it is considering are code-for-letters, or really numbers, and this is accomplished by a prior instruction: "Consider the next N numbers as code."

So when you see flip comments typed out on a printer by a computer, e.g., "That was a stupid mistake, kid," unfortunately you have to imagine the entire sentence, spaces and all, stored in coded form inside the machine, by a flip program, written by a flip programmer.

4. COMPUTERS CAN MAKE DECISIONS

The flip sentence was not remarkable, but the decision to print the sentence was made by the computer and so might be remarkable. Let us consider some computer decisions, as typed out on the printer.

A. "Mount new reel of magnetic tape on Tape unit #3." How taken? A possible method is this: a count of 36-digit binary numbers being written on the tape passes the 2000 mark. Mysterious? Hardly. Remarkable? Hardly.

B. "Kt to K4." This is chess jargon for moving a knight to the 4th square of the column in which the king was located at the beginning of the game. This decision is remarkable. It can be based on a consideration by the computer of all possible moves by it and by its opponent for several turns ahead. The trick then is to attach some kind of value number to each final position which is considered and to maximize this number. Here is a good example of my general remark that in applying computers to human problems you really study the problem. But now a further step can be taken. One type of value assignment can be played against another by the computer and it can select the one which wins more frequently. The computer *learns*. This process is an example from that pioneer branch of the computer applications field called "self-organizing systems." Evidently the capacity for self-organization is one of the highest human faculties, so there is strong interaction here with the development of human values.

5. TWO PROBLEMS AT THE THRESHOLD OF COMPUTER CAPABILITY

A. The Spoken Language

First, computer speech. The computer in speaking could use a prerecorded vocabulary of words spoken by a human. Let us ignore this slightly trivial possibility. Then the computer must compose words out of the grunts, buzzes, and whistles available to humans. This can be done, and again we learn much about the mechanism of human speech. However, the pronunciation, especially in English, is not contained in the written word, and the sentence rhythm and inflection contain an enormous amount of information. Consider "Oh" (skeptical) and "Oh" (resigned). Because speaking computers are rather a luxury than a necessity, the problem is not solved now in practice, as far as I know.

Next, computer listening. Computers can understand one person's voice for quite a range of vocabulary, and also a range of voices for a limited vocabulary (such as the names of the numbers). Since a voice-operated typewriter would be of economic importance, there is considerable work going on here, but I believe that the problem is not yet adequately solved. It will be amusing if some day we find a voice typewriter with a dial on the side which can be set to "Southern Drawl," "New Yorkese," "Texas," or "English English." Our French-speaking colleagues can use the machine as an objective test of their accent, with a choice of which possible English they would like to try for.

B. Painting and Musical Composition

For painting, the computer has an interesting medium, the color TV screen. Don't judge by the usual crude home color receiver, but instead look sometime at the $4000 receiver wired to the camera in a color studio. The saturation of the colors, to use a technical term, exceeds that from any painted surface. The resolution and draftsmanship is as good as most human painters can manage. The texture isn't very interesting and the computer should work on that.

In writing music, the computer can use recorded instrument tones or mix up its own sounds as in the new "electronic" music.

At this point I tend to feel a value judgment coming on. I tend to say music and painting are for individual human enjoyment. Let the computer produce any set of paintings it wishes, and I will preserve one if I like it and even hang it up on some wall. Let someone even ingeniously program some painting rules into the computer and, again, if anyone likes the result, let him keep it. Best of all, for someone who likes to paint but simply can't draw a straight line, let him use the computer to make himself some paintings to hang on his own wall. But do not let anybody ascribe a value number to a painting, human- or machine-made, and try to maximize that number. What I am really suggesting, I guess, is that one permanent human value should be a clear master-slave relationship between humans and computers.

6. ANALYSIS OF BUBBLE CHAMBER PICTURES

Now having ranged over what must be something like the gamut of computer capabilities, let us settle down and consider a computer in relation to the problem of bubble chamber data analysis.

First let us look at the typical bubble chamber picture, Figure 4, a photograph from the 30-cm liquid hydrogen bubble chamber of the CERN bubble chamber group. Being typical, the picture is not remarkable from the strange-particle physics point of view. As a piece of visual information we note some general characteristics: (1) The picture is mostly "empty." (2) Bubbles not belonging to tracks are a small percentage of all bubbles. (3) The precision with which bubble centers lie along a particle track is $\approx 1/30$ bubble diameter or less.

Now some logistics. Several million exposures of a chamber can be made per year. About a hundred thousand exposures are required, in a typical experiment, to obtain a sample of a few tens or a few hundreds of strange-particle events from which the behavior of known particles and the existence of new ones can be deduced. For very rare events, just scanning the pictures is the bottleneck in doing the experiment. Very quickly, however, as one learns under what circumstances a new particle is favorably produced, the bottleneck shifts to measurement. The principal means for identification of particles is balancing momentum and energy at the collision where they are produced or afterwards at points where they decay. A particle's momentum is measured by the deviation of its track from a straight line in the chamber's magnetic field, and at high energy this deviation may be only a few bubble diameters. Therefore very high precision measurements are required, and, even with the great ingenuity shown in the design of present measuring machines, the process is excruciatingly slow. Finally, for experiments in which several thousand events of one type are analyzed, say for the dependence of the process on energy and angle, the labor just in bookkeeping and in following up dubious cases can be prodigious.

Just two years ago Brian Powell and I found ourselves at CERN, newly enough arrived that we had not accumulated a set of detailed responsibilities. Powell had done cloud chamber work at the Pic du Midi in the Pyrenees and was instrumental in setting up the Imperial College measuring machine for bubble chamber pictures at London. I had used nuclear emulsions in low energy nuclear physics and had worked out some electronic track detection techniques for emulsions. We became interested in trying to improve existing bubble chamber techniques and were encouraged in this endeavor by Yves Goldschmidt-Clermont, who promptly organized a Crazy Ideas seminar. At the time we did not really realize that the measuring process is so often a bottleneck in bubble chamber physics, but we tackled this process rather than scanning because it seemed much more possible. At the same time we wanted any new system to be open in the direction of automatic scanning, as a future development. We wanted high speed, but were all the same not interested in any ideas which would not compete favorably with existing measuring machines from the point of view of accuracy.

The idea perhaps most essential to producing the new system was this: rather than using human guidance in series to lead a measuring machine along the tracks of an event to be measured, find a way to measure routinely every bubble in the picture and then (somehow) apply a mask over this undigested mass of data to select only the bubbles belonging to the event. Human guidance in parallel, so to speak. At first the mask idea was taken very literally indeed and we imagined painting colored transparent bands called "roads" over the interesting tracks. Then if a spot of light carried out a regular rectangular TV-like scan over the entire picture, the instrument would know to digitize only bubbles lying in the roads. The

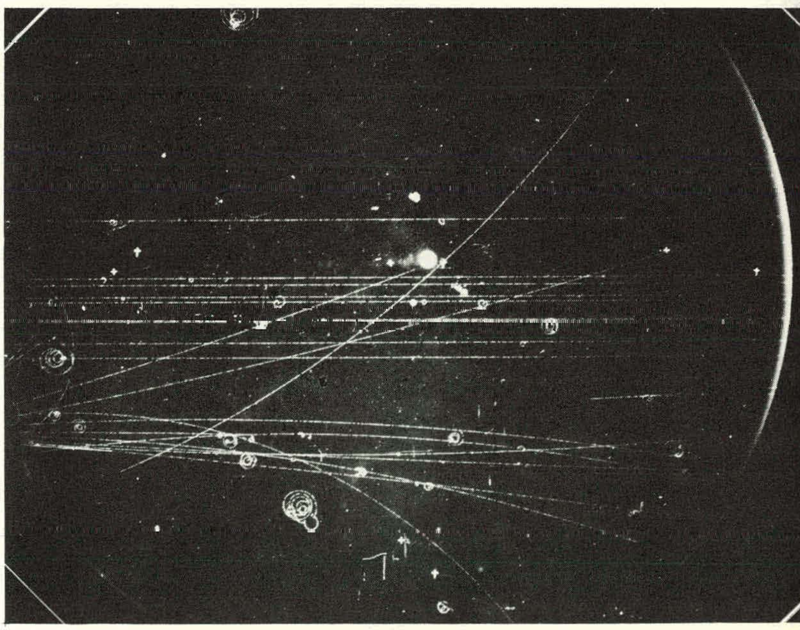

Figure 4. A typical bubble chamber picture.

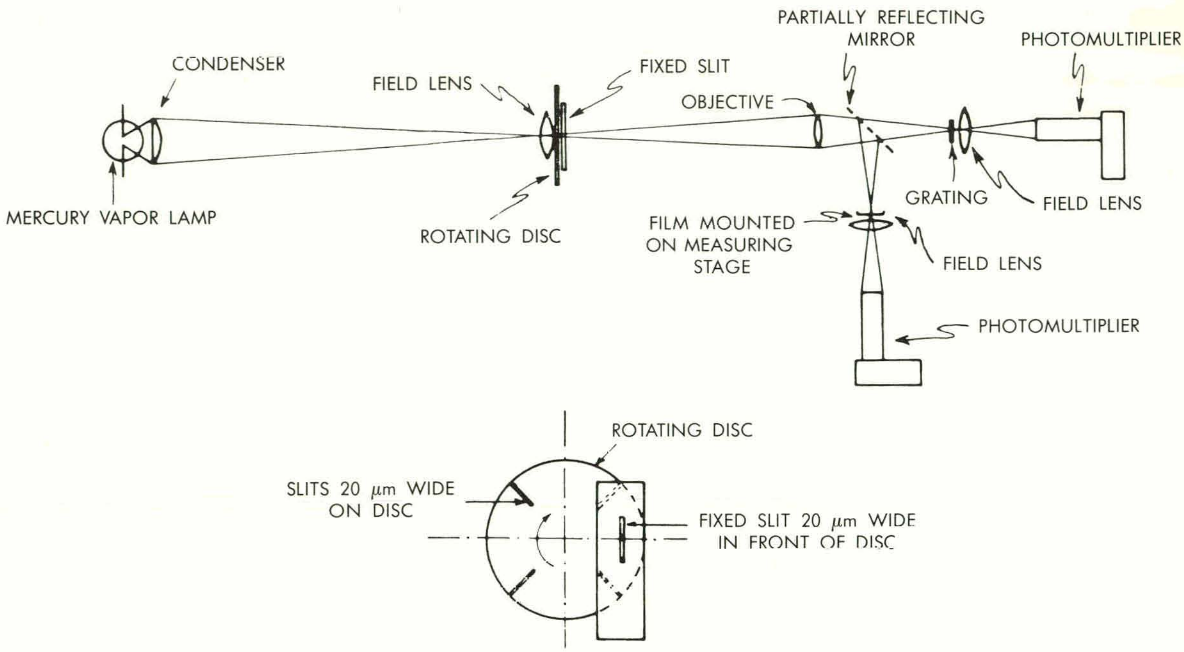

Figure 5. Scheme for generating the flying spot.

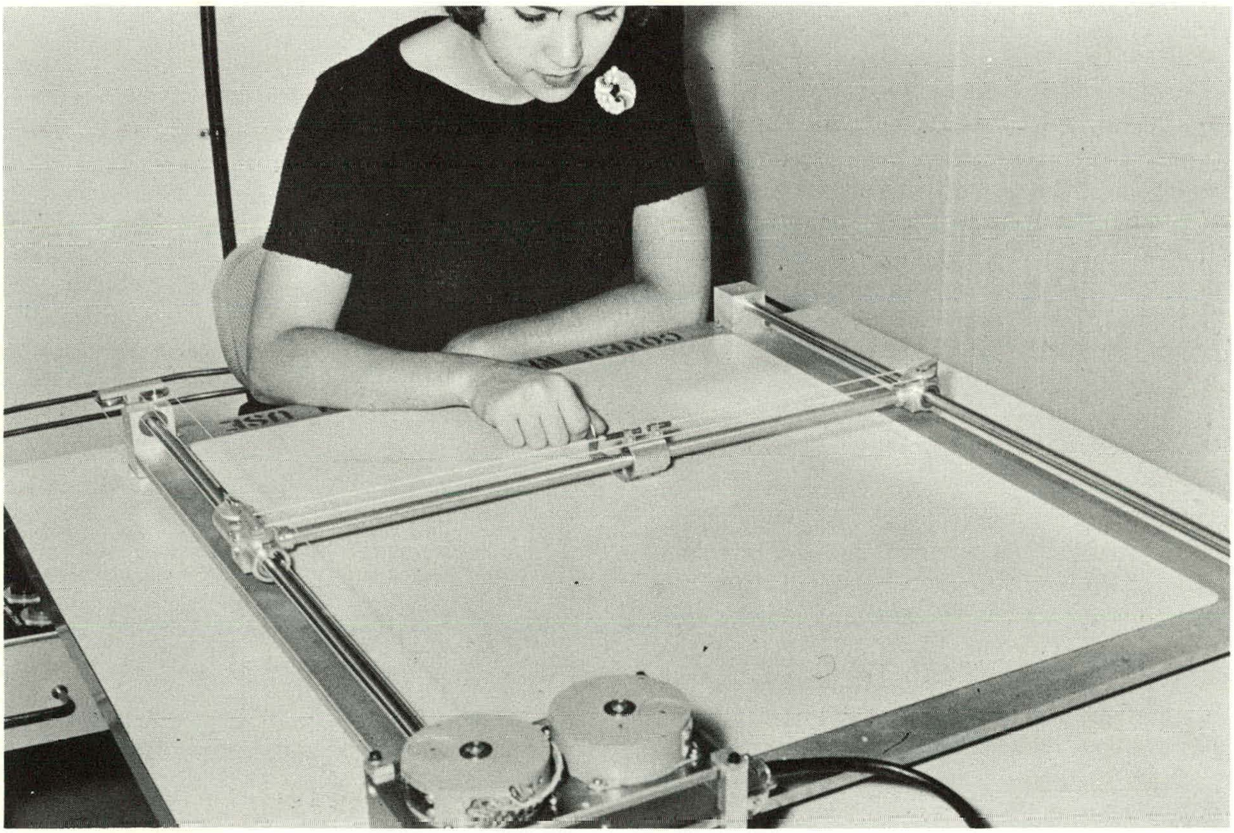

Figure 6. A simple rough digitizer accurate to about ±½ mm over an area 60 cm².

selected coordinates were to go onto magnetic tape for later transfer to a computer. The CERN measuring machines are called IEP's, IEP standing for "Instrument for Evaluation of Photographs," and to IEP's 1 to 5 we added an IEP-X.

The first job was to try out the flying spot of light which scans the photograph. In principle the flying spot could have been a flying line, but a line at arbitrary angle seemed to require an oscilloscope tube for generation, and the fineness of line available from such tubes seemed then (and still seems) marginal. Further, the time to sweep over the entire negative with a line segment at the many angles required seemed too long. Finally, it is difficult to determine the end points of a track with a line detector. So we chose a flying spot. Of course, then, if the computer is to understand the picture unaided, it must reconstruct lines from points, so there is still plenty of room for competition between the methods.

It proved possible to produce a spot $\approx 15\ \mu$ (or half a thousandth of an inch) in diameter by mechanical means familiar to early workers in television. As indicated in Figure 5, radial slits about one thousandth of an inch wide are mounted on a disk which turns at 60 rev/sec. About ¼ mm away from the moving slits is a fixed slit of the same width. Light from a mercury arc passes through the intersection of the two slits. The resulting spot is imaged at a 2:1 reduction onto the film. As the disk turns, the spot traverses a line on the film. We found that the light transmitted through the film was attenuated by 25 to 50% as the spot crossed a bubble image. With a little advice from R.L. Garwin of IBM and Columbia, who was also at CERN that year, we were able to find the center of the bell-shaped attenuation curve to about 1/10th of its width, and this gave us, in principle at least, a measuring accuracy about equal to that of the best existing measuring machines. There remained the problem of knowing where the spot was on the film when the center of a bubble image was detected. This was accomplished by splitting the light from the spot generator with a partially silvered mirror and forming a second spot image on a precision grating. By counting grating lines and interpolating between them we were able to fix the position of the spot at all times during the sweep to an accuracy of about 2 μ. The film itself is mounted on a precision stage whose position is digitized by standard methods. As the stage is driven along at (typically) 1 in/sec the flying spot sweeps across the film in lines separated (typically) by 1/20 mm (or 50 μ).

The basic flying spot studies were completed at CERN a year ago last fall. On my return to this country our system ideas were modified and developed considerably by the criticism and suggestions of American physicists, especially Alan Thorndike. Dr. Thorndike rather objected to the large amount of magnetic tape handling which IEP-X would have required and advised us of the existence of a mechanism for rapid transfer to internal memory available on IBM 709/7090 machines called the "Direct Data Connection." It soon became clear that we could read all the coordinates from a picture directly into the computer if we wished and not just the coordinates lying in roads. This was fortunate, because the making of literal roads had been giving considerable trouble. Brian Powell and Dr. Ian Skillikorn, now at this Laboratory, had succeeded in making rather beautiful road maps in the form of masks on separate pieces of film prepared at the scan table. However the extra bookkeeping involved in handling the masks and the problem of aligning picture and mask at the flying spot digitizer caused a quite general lack of enthusiasm. We then revived an idea proposed by Ross McLeod of CERN during the original Crazy Ideas seminar. This was to make a few rough measurements at the scan table, transmit these measurements on an IBM card to the computer, and let it draw curves through the points which, with error bands on each side, would then constitute the roads. McLeod had proposed having the computer put out road coordinates which would be used to select coordinates outside the computer. The system finally adopted however was to allow all coordinates to enter the computer and road-select inside. The principal reason for our choice was that the system is then one stage closer to unassisted pattern recognition by the computer. The rough measurements at the scan table can, we now know, be made conveniently by several techniques. Figure 6 shows a rough digitizer designed at Brookhaven by Carl Goodzeit and built by Gottfried Szongott. Rose Ann Giambalvo is posing for the picture. Our current idea of the ultimate in rough digitizers is one under development by Marty Rosenblum and Bob Chase of Instrumentation, following a proposal of Teager at MIT. It is being

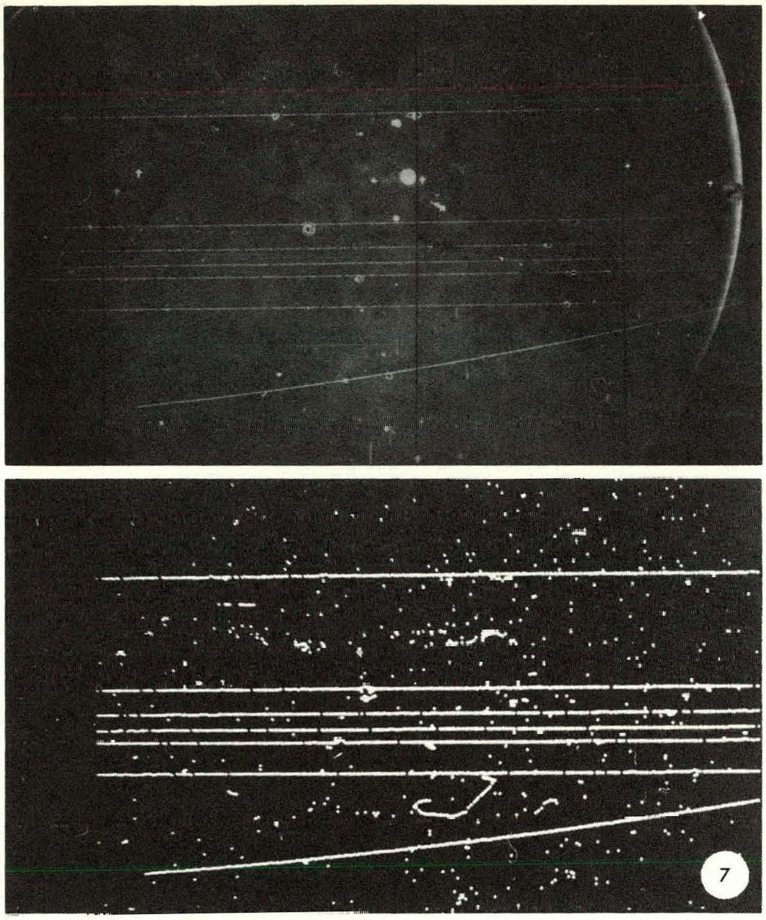

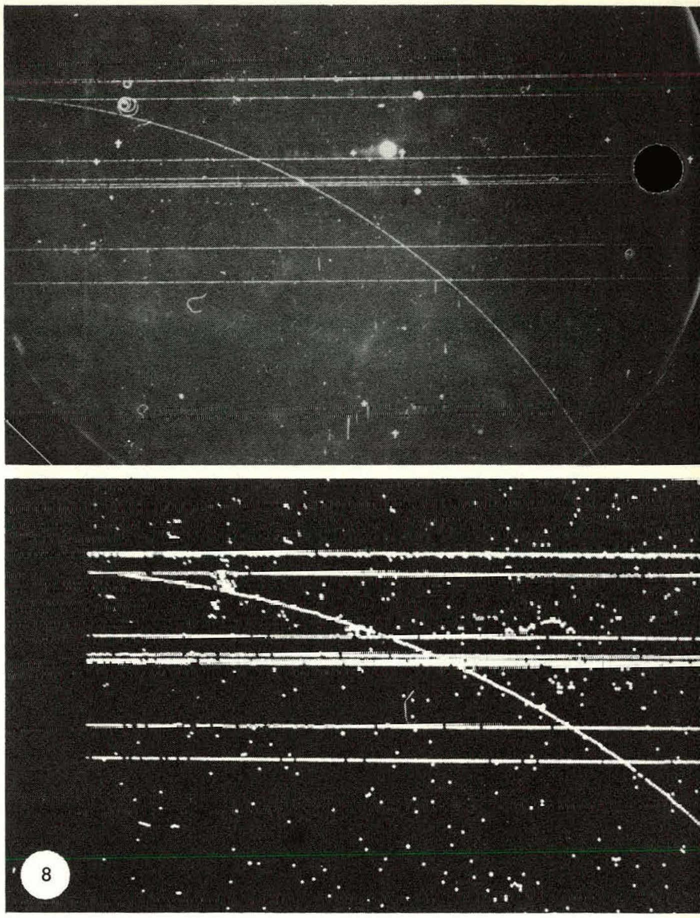

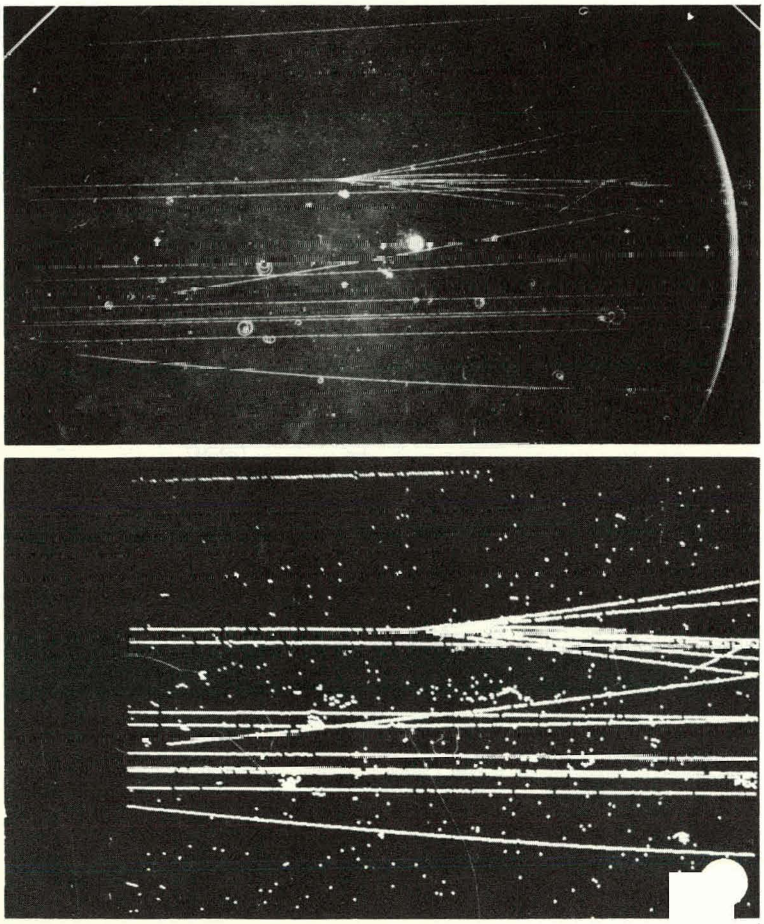

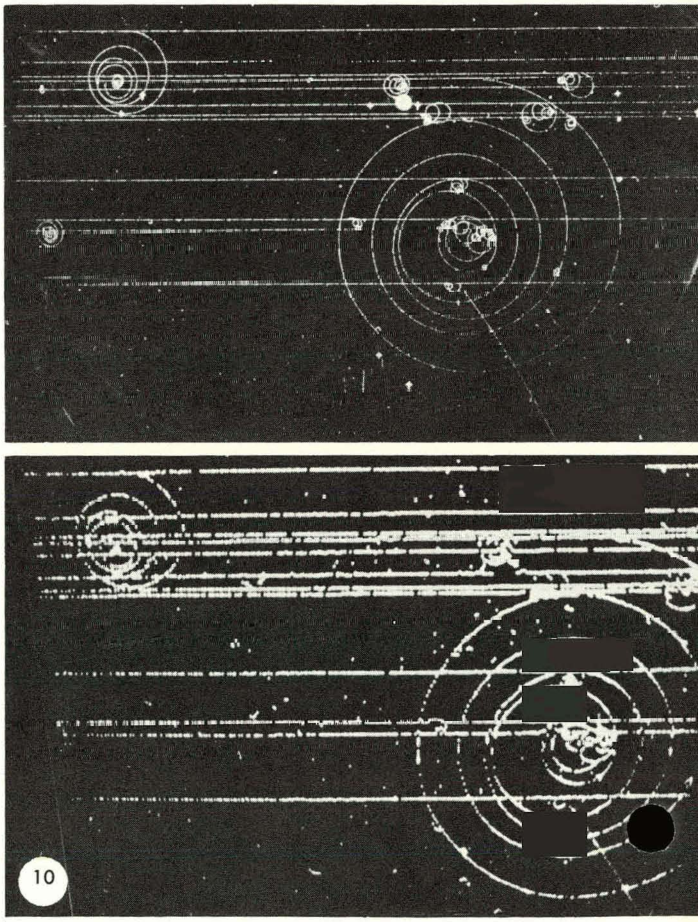

Figures 7 to 10. Comparison of original photographs with the results of measurements by the prototype flying spot digitzcr at CERN.

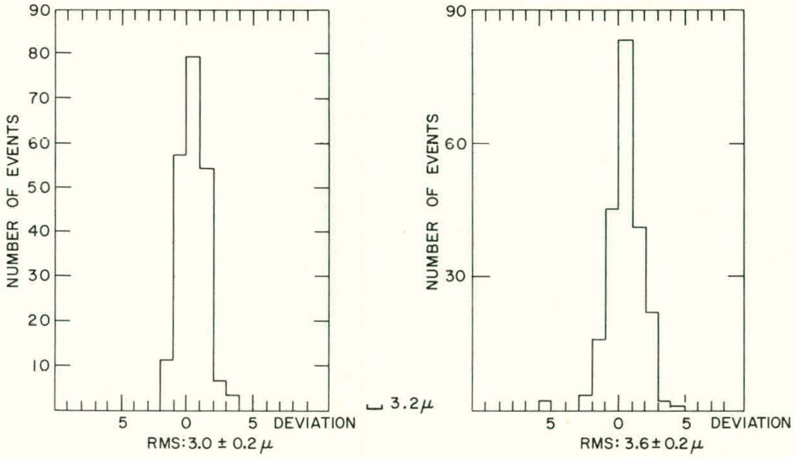

Figure 11. Histograms showing the deviation of flying spot digitizer measurements from a smooth curve.

constructed by Clarence Porter. There is interest in their instrument from a number of quarters, and I will leave it to them to describe it another time.

A year ago our data analysis system was conceptually complete, and it was decided to try it out with the CERN prototype flying spot digitizer. A Direct Data Connection was acquired by Dr. Kowarski, head of the Data Handling Group, for the CERN 709. Electronic circuitry to permit the flow of coordinates from film to computer was designed here by Tor Lingjaerde and myself and built partly here and partly at CERN. Howard White of Berkeley provided programs to handle the coordinates inside the 709. Jerry Russell, then of Berkeley, now at Brookhaven, got the 709 Direct Data Connection operating properly and generally helped out with electronics. Not many days after the necessary hardware was physically there, about the end of last May, coordinates began flowing into the 709 at the rate of several thousand a second and tens of thousands per picture. The coordinates were printed out for analysis by hand, and also shipped on magnetic tape to Berkeley where a 709 was equipped with a display oscilloscope. Some of these pictures, repainted by the computer, are shown in Figures 7 to 10. Although the display oscilloscope has much lower resolution and precision than the data, it is useful for showing what is digitized. The precision of the measuring machine is shown in Figure 11. The histograms show the scatter of individual points about a smooth curve. The flying spot digitizer makes many more measurements per track than conventional machines, so that it is possible to average 20 to 30 neighboring points to get master points. These master points have been shown by Robert Palmer of Brookhaven, working with the prototype at CERN, to deviate from a smooth curve by appreciably less than 1 μ, and this precision is nearly an order of magnitude higher than that available from conventional machines.

We were all quite satisfied with these results and returned to our respective Laboratories to attack the problems of actually getting complete analysis systems on the air. David Lord and John Burren of Harwell, who had participated in the CERN tests, received approval to go ahead with a system for the British National Institute for High Energy Research. The Wilson Powell group at Berkeley, especially Howard White, Jerry Russell, and Jack Franck, began operations. A small group of us here, including initially Alan Thorndike, Ronnie Rau, Bob Palmer, Ed Rogers, Neil Webre, and myself, started work. At MIT, a group working under Professor Irwin Pless began a somewhat related development. In particular, Pless is pioneering in the use of a precision oscilloscope tube as spot generator. These four Laboratories and CERN are collaborating closely, exchanging results and occasionally people. One very recent result from CERN obtained by Brian Powell and M. Benot may well be of great usefulness. They have shown that the flying spot digitizer can be run so as to give the density of bubbles along a track with very nearly the same precision as the most detailed hand measurements. The bubble density gives the speed of the particle and therefore (in combination with track curvature) its mass, at least in favorable cases. Until now it has often been necessary to ignore bubble density information because of the labor of acquisition.

At Brookhaven, the basic problems now being reduced fall into four categories:

The first is proper preparation of the film at the bubble chamber. Ed Hart, Bob Louttit, Marty Rosenblum, Karl Abrams, and Russ Willoughby have been preparing proper fiducial marks and turning a perfectly clear picture number into a confusing code of black and white squares that machines understand better. This work is now nearly finished.

The second is equipping scanning tables with rough digitizers and arranging for transmission of road information to punched cards and thence to the computer. Jack Cockrill and Carl Goodzeit have been largely responsible for this work, which is well along.

The third is the construction of the mechanics and electronics of the flying spot digitizer itself. Carl Goodzeit, John Leavitt, Carl Pozgay, Joe Scheliga, John Tagliavia, and I have been doing mechanics, and Ed Rogers, Jerry Russell, and Ed Frantz electronics.

The fourth is the important and difficult area of programming. Alan Thorndike has been coordinating an effort which includes Neil Webre, working until now at Berkeley, Phil Conolly, Jerry Friedman, and Joyce Tichler.

Figure 12 shows our rotating disk, here stationary, with eight slits (too fine to see) contained in these pockets. Professor E.M. McMillan, director of the Radiation Lab at Berkeley, has shown us how to curve these slits to give a constant speed of scan. Figure 13 shows the over-all machine in assembly. The precision stage was designed by Jack Franck. A high speed film transport, about half built, will occupy the open space to one side of the stage.

Finally, we have a small group consisting of Bob Marr, Bill Beard, and George Rabinowitz, who are working on programming the 7090 computer to understand the coordinates fed it by the flying spot digitizer directly, without roads and in fact without any human guidance at all. Marr and Beard are following up methods introduced by Professor John Pasta of Illinois when he was here

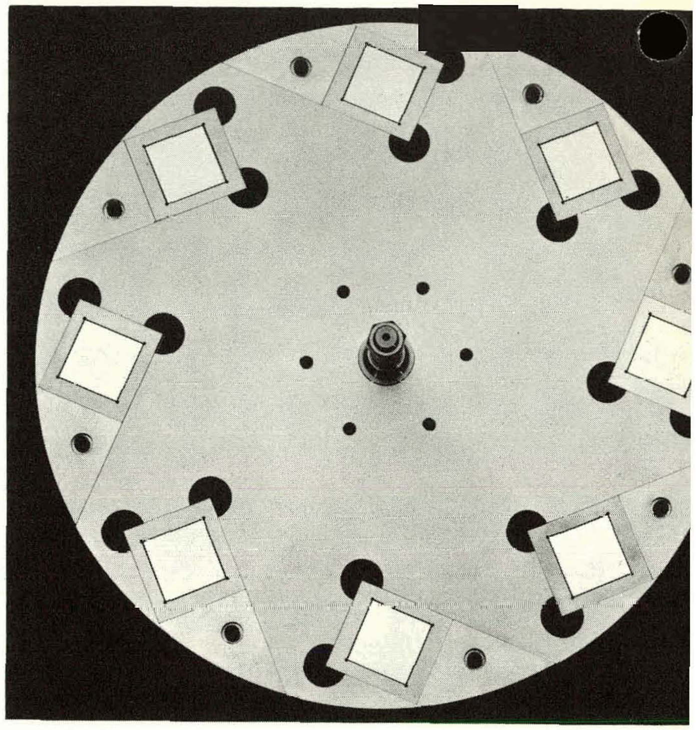

Figure 12. The rotating disk of the Brookhaven flying spot digiti

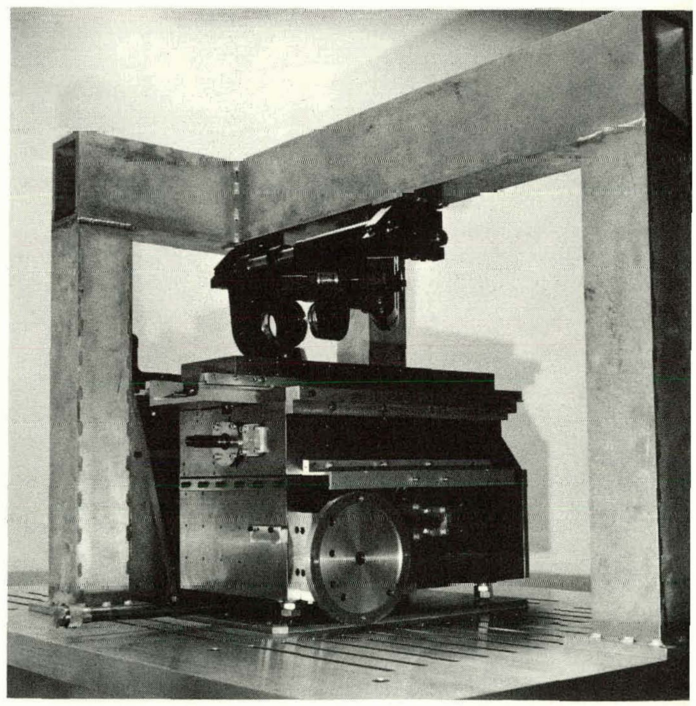

Figure 13. The Brookhaven flying spot digitizer in early stages of assembly.

Figure 14. The first results of a pattern recognition program of George Rabinowitz. The line segments plotted have been reconstructed from individual unlinked points.

last summer. Their 7090 program is very nearly ready to go and has already given interesting results in a trial version run last fall on Merlin. Rabinowitz has devised an independent attack on the problem which ran for the first time on a complete picture last week. The raw materials for both these recognition programs are magnetic tapes containing all the bubble coordinates for typical Brookhaven 20-in. chamber photographs, digitized last December by Powell and Palmer using the CERN prototype flying spot digitizer.

Rabinowitz's first results are shown in Figure 14. TheRabinowitz program uses essentially only the criterion of "nearness" to link up isolated coordinates into track segments. The input to the program is the set of bubble coordinates as they are measured, line by line. The output is separate tables of coordinates, each table believed by the computer to be a piece of a track. What we show here is a plot of the linked-up pieces of track.

The 7090 computer takes about 5 sec to achieve this degree of understanding of each of the separate stereo photographs taken of the Brookhaven 20-in. liquid hydrogen bubble chamber. Two factor-of-two increases in speed are on the horizon, and, if they can be realized, a computer will be able to perform this very imperfect scanning on about 1000 stereo triads of photographs per hour. More important, perhaps, the direction of a track element is to be used in selecting likely candidates for points which extend a known length of track, and this criterion is appreciably more powerful than simple "nearness." On questions of over-all geometry, e.g., whether three tracks have a common vertex, nothing has so far been attempted.

7. SUMMARIZING REFLECTIONS

These general activities amount in some sense to enabling a computer to see. In the course of the work we certainly are gaining a very fair appreciation of the nature of the problem of recognizing track patterns. Perhaps, when the computer gets tolerably competent at it, we can turn our attention to another problem – how it is that humans handle this problem with ease and so astonishingly well.

LIST OF BROOKHAVEN LECTURES

1. Radioastronomy and Communication Through Space, BNL 658
Edward M. Purcell, Physics Department — November 16, 1960

2. Current Ideas on the Endocrine Regulation of Cellular Processes, BNL 685
Irving Schwartz, Medical Department — December 14, 1960

3. Inside the Protein Molecule, BNL 649
Werner Hirs, Biology Department — January 11, 1961

4. Nuclear Chemistry Research With the Cosmotron
Gerhart Friedlander, Chemistry Department — February 15, 1961

5. Neutron Physics Of and With the High Flux Beam Research Reactor, BNL 664
Herbert Kouts, Nuclear Engineering Department — March 15, 1961

6. High Energy Accelerators
Ernest Courant, Accelerator Department — April 12, 1961

7. Dislocations in Crystal Lattices
George H. Vineyard, Physics Department — May 17, 1961

8. The History of Cosmic Rays in Meteorites
Oliver A. Schaeffer, Chemistry Department — June 14, 1961

9. The Physics of Semiconductor Radiation Detectors, BNL 699
G.L. Miller, Instrumentation and Health Physics Department — September 27, 1961

10. Theory of the Gene
Milislav Demerec, Biology Department — October 18, 1961

11. Physics of Fundamental Particles
Maurice Goldhaber, Director, Brookhaven National Laboratory — November 15, 1961

12. Excessive Salt Intake and Hypertension: A Dietary and Genetic Interplay
Lewis K. Dahl, Medical Department — December 13, 1961

13. Galaxies, BNL 710
Otto Struve, National Radio Astronomy Observatory — January 17, 1962

14. A Computer Learns To See, BNL 725
Paul Hough, Physics Department — February 14, 1962