6 October 1967, Volume 158, Number 3797

SCIENCE

Automation in the Laboratory

On-line computers are providing new freedom in the design and conduct of experiments.

R. J. Spinrad

One way to gauge the nature of current research is to observe the skills in which the graduate student is trained. For, although the scientific method is immutable, the *techniques* of science change continually. In the 1930's the student was expected to be a proficient glassblower. Electronics was the *sine qua non* of research in the 1940's and 1950's.

Today the computer is king. The science student, in his undergraduate years and through all his subsequent training, is schooled in the art of automatic computation. But, beyond its classical role in scientific calculation, the computer is now being used in the laboratory as an integral part of the experimental apparatus itself. This blossoming field of laboratory "automation" is the subject of this article (*1*).

Levels of Automation

First, let me discuss the steps involved in going from a classical experimental situation to a fully automated one.

It is possible to talk in quite general terms about an experiment and its "input" and "output" (for physiology, read "stimulus" and "response"; for neutron-diffraction analysis, read "position crystal" and "read counter") (see Fig. 1). I am concerned here not with the nature of the experiment but with the technique.

At the simplest level, the input is a manual operation and the output—noted from some meter or counter—is entered, by hand, in a log book. Or the results, perhaps as a function of time or of some input parameter, are plotted on a strip chart or *x-y* recorder.

The first step toward "closing the loop" for the automated laboratory occurs when the experimenter decides to generate his data in machine-readable form. Appropriate analog-to-digital converters are incorporated in the output apparatus, and the results appear directly on punched paper tape or cards or on computer-compatible magnetic tape. (Often a printed record is provided simultaneously so that the experimenter can maintain close contact with his experiment.) The output data are computer-processed later. The results of the analysis are used by the experimenter to determine the next steps in his experimental program.

Often, however, the judgment of the scientist is required only a few times in the course of the experiment. Wording it differently, we might say that there are subsets of the input series (a simple succession of stimuli, of settings, of positionings, and so on) which are so frequently used that they may just as well be made automatic. This automation may take the form of a motor slowly turning a potentiometer or of a sample changer progressively exposing different items. Or, to state this in general terms, the input sequence could be arranged to be "programmed" by a series of codes punched in, say, a paper tape. The digital codes of the paper tape would then be converted into the analog signals necessary to "drive" the experiment.

The experiment is then "digitized." A series of digital input commands directs it, and the results are made available in digital output form. Figure 2 represents this situation.

Plan of the Experiment

There must, of course, be some *plan* for the experiment, either carefully formulated or existing only as a series of notions in the mind of the experimenter.

In the steps discussed above, the plan is evident in at least three places: (i) the generation of the digital input series, (ii) the computer program that processes the digital output data, and (iii) the scientific judgment that goes into linking (i) and (ii).

In the simplest cases, where the codes are compact, the input commands are punched by the experimenter on a tape-preparation typewriter. More often, however, the experimental sequence is generated by a separate computer program which details all the elementary steps necessary for one run (or scan, or session, or trial). The input tape for the experiment is produced directly by the computer as a part of the program's operation.

Let me review the partially automated experiment up to this point. The experimenter prepares a series of "macrocommands" describing one run. These commands are entered into a computer, along with the appropriate generating program, and are there elaborated into the detailed series of "microcommands" necessary to drive the apparatus. The paper tape resulting from this process is fed into the experimental equipment, which then, a step at a time, executes the run. As the run proceeds, the output data are being generated on another paper tape.

The author heads the Computer Systems Group at Brookhaven National Laboratory, Upton, New York.

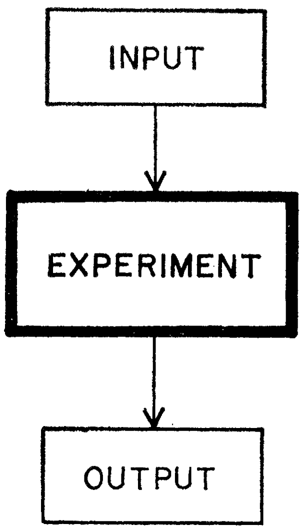

Fig. 1. Generalized experimental configuration.

At the completion of the run this output tape is fed into a computer, along with the data analysis program. The printed (or plotted) results are presented to the experimenter for his consideration.

On the basis of what he learns, he prepares a new program of macrocommands to execute a different run. The process is repeated until the experimenter is satisfied.

Closing the Loop

It requires no great flight of the imagination to see how to "close the loop." Clearly, if the scientist can formalize the mental processes he employs in going from results to input commands, those steps can be programmed. If they can be programmed, the data analysis program and the input preparation program can be joined with this "decision-making" program to form a single master program. The master program would accept the data from run i, process them, apply the preprogrammed decision criteria, and, as a result, generate the driving tape for run $i + l$.

This is not as fanciful as it may seem. There are a number of laboratories working in just this way today. Of course, complete experiments of such tractability are in the minority. However, even the most sophisticated experiments have parts that are susceptible to this kind of treatment. And, as understanding grows, longer and longer portions can be systematized and codified.

In addition, the computer can aid the experimenter in many other ways, even when the situation is so complex that he must remain in active control.

"Going On-line"

Conceptually, the introduction of the computer program to take the experimental cycle from output back to new input completes the picture of the automated experiment—except for the important element of timeliness. The last, and most crucial, step is that of radically speeding up the whole process. This is done by "going on-line" (that is, by electrically connecting the experimental apparatus to the computer).

The most inefficient links in the experimental loop described so far are those intersected by the vertical, dotted boundary in Fig. 2. At the right in Fig. 2, the experimental process proceeds at a rate limited only by the nature of the apparatus and by the experimental materials and conditions. The rate at the left of Fig. 2, representing the programmatic execution of the experiment plan, is limited only by the microsecond speeds of the digital computer. The bottleneck exists in the delays incurred in generating the output paper tape, transporting it to the remote computer, waiting for "computer time," and then carrying the new input tape back to the experimental apparatus—which has been idle in the interval.

The obvious final step is to do away with the paper tape and directly connect the experiment output to the computer input and the computer output to the experiment input. Figure 3 represents the final system.

The effects of putting the computer on-line to the experiment are much greater than would be immediately assumed. Because of the immediacy of the interaction, it is no longer important to "batch" the experimental measurements into "runs." Detailed numerical analysis of each datum is as valid a part of the experimental plan as group analyses on the results of a day's measurements.

This has the effect of freeing the scientist from the constraints of a kind of 19th-century experimental orthodoxy. Experimental designs become more imaginative and bolder. Data can be taken, for example, just to the point where the desired statistical accuracy is achieved—without the need for conservative overrun. Each measurement can be more to the point because all the preceding data—reduced and analyzed "live"—show the way. In fact, for the first time the scientist can "program" his entire experiment, mixing, as necessary, measurements, delays, computations, decisions, reporting functions, and requests for guidance. In actual experience, the changes in the nature of research are striking.

The Apparatus

What of the equipment necessary for the automated laboratory? How does it differ from that already in use?

Of course, the prime addition is the computer. Generally speaking, laboratory computers fall in the small-to-medium size range. Depending on the scope and complexity of the experimental environment, the computer configuration will cost anywhere from $10,000 to $300,000. (That is quite a range, but so also is the performance range of these systems.)

I will not discuss the selection of computer equipment and programming systems here, except to say that it is perhaps the most complex purely technological issue the experimenter is ever likely to face. Fortunately, more and more of the computer manufacturers are becoming familiar with the "laboratory" market and are ready to provide expert (though possibly biased) counsel.

At the other part of the loop, the experimental apparatus must differ from the conventional in that all of its outputs and inputs must be digitized. This is a broader statement than may at first seem to be the case. It, of course, implies that the readings of scalers, rate-meters, photometers, and so on must be in coded form. But, more than that, it implies that start-stop switches, gain controls, prescaler settings, mode switches, and the like must be capable of being operated by electronic signals.

Increasingly, manufactured instruments are being provided with these features. In many instances, however, they must be added to, or built into, existing equipment.

Connecting the digitized laboratory apparatus to the digital computer is a straightforward engineering job which is, these days, made easier by the ready availability of modular electronic "building blocks." These "logic cards" come in great variety from many manufacturers (including some computer manufacturers). The rules for their inter-

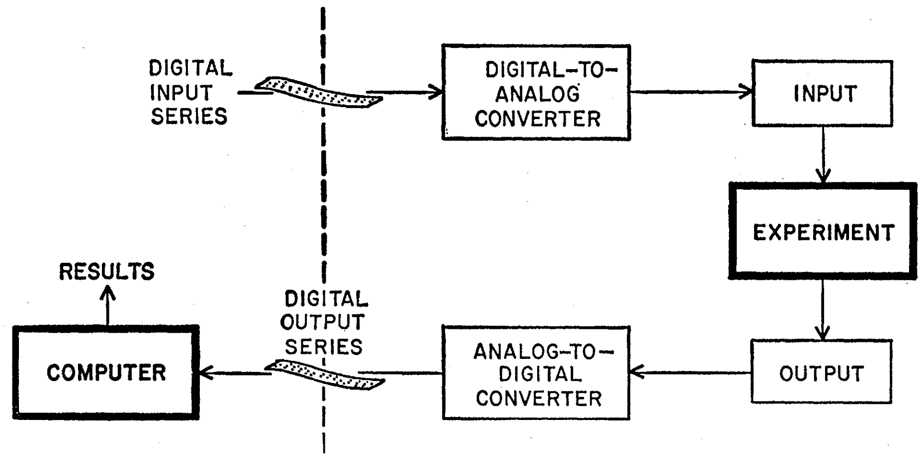

Fig. 2. "Digitized" experiment run by punched paper tape. The raw data are punched into the paper tape.

connection are simple and precise. Little experience or training is required to establish a perfectly satisfactory computer-experiment "interface."

In the jargon of the field, we can describe the on-line computer as having "front-end" and "back-end" attachments. The experimental apparatus just discussed is the front end. At the back end are a variety of devices which are used to provide flexible and effective control of the experiment.

Without question, the most ubiquitous of these devices is the on-line typewriter. The typewriter is used to issue commands, to receive reports or comments, to enter control parameters, to log results, and, in general, to provide for any communication function not explicitly handled by other apparatus.

In the simpler systems it is the only communication device and, as such, "talks" for the equipment in the "dialog" between the experimenter and the experiment.

An interesting by-product of the use of the typewriter as both a command and an output medium is the fact that the typewritten sheets are used in lieu of a handwritten experiment log. Through proper programming the record can be made to show time and identification for every entry or "transaction."

All but the simplest systems include a magnetic tape unit as an item of peripheral apparatus. If the experimenter has followed the prudent course of selecting a computer-industry standard tape unit, he can use the tape as a medium of interchange for both data and programs. Furthermore, he will be readily able to process his data on other computers if the need arises.

Experience has shown that one of the most useful back-end devices in the laboratory environment is the cathode-ray-tube display. With it the experimenter can display graphs, scatter diagrams, numbers, or textual material. The display can be "live," showing raw, partially processed, or fully processed results.

Greatly enhancing the usefulness of the cathode-ray-tube display are various "pointing" devices (for example, the "light pen" or "trackball") which enable the experimenter to designate to the computer system a particular item of interest on the current picture. He might want to single out anomalous data points for special analysis. The pointer could be used to indicate limits between which integration under peaks is to take place, or to sketch in a baseline to be used for subtraction of background.

Sometimes the display is used to portray, in graphical form, the progress of the experiment. Often it is used instead of the typewriter to present a rich body of data on the basis of which the experimenter must choose among alternative courses. It is, I believe, the single most flexible communication tool.

Other back-end devices abound. The system can use printers, *x-y* plotters, aural devices, and so on. Appropriateness is the only criterion.

System Software

It is both conventional and convenient to separate "hardware" and "software" when discussing computer applications. This separation, however, often masks a most important characteristic of effective computer systems design—the fact that the "systems" part of the design demands the simultaneous, interwoven development of both hardware and software aspects.

Invariably considerable savings (in both time and money) can be achieved by freeing the mind of preconceptions concerning what a software function is and what a hardware function is. It may be, for example, that you do not have to buy motor speed controls because the on-line computer can handle the task programmatically. Banks

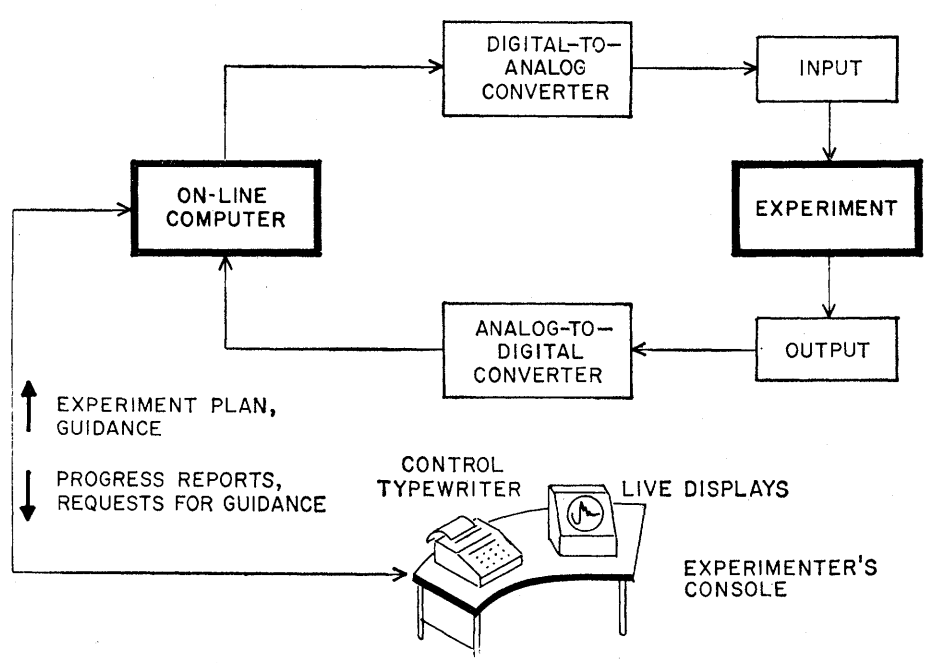

Fig. 3. "Automated" experiment. The on-line computer executes the experiment plan; the experimenter exercises only supervisory control.

of binary counters may be less expensive than their decimal equivalents, and then readout by way of a binary-to-decimal conversion routine may be performed by the computer. Or perhaps the counters are not needed at all; the computer can be programmed to act as a set of counters.

On the other side, certain kinds of complex signal processing functions may still be most expeditiously handled by the hardware. Combining data into the proper arrangement for entry into the computer is probably best handled by a patchboard. Even some of the rather sophisticated control programs can be considerably simplified through the use of specialized hardware "priority interrupts."

A skilled system designer is probably the "investment" with the best payoff for the laboratory that is thinking of "going automatic." Smaller laboratories that cannot support a full-time specialist still have many sources to draw upon. Workers in university laboratories should be able to interest some of their engineering or computer-science colleagues in their automation problems. For nonaffiliated laboratories, there are many expert "systems houses" and consulting firms. Some—though not many—of the computer manufacturers have competent advisory services.

The Experimenter's Programs

The well-designed system separates —or perhaps I should say "insulates" —the experimenter from the red tape and organizational details of the system control program. The user sees available to him the macrocommands he needs to manipulate his experiment and a programming language in which to express the computational and decision-making parts of his experimental plan.

In the earliest days of on-line computer systems, the programming language was always machine language or a symbolic assembler language. The last few years have seen the introduction of higher-level (compiler) languages such as Fortran or Algol. This is a welcome development because it opens up the field of laboratory automation to a much wider audience—those trained in Fortran or Algol. In addition, it makes available a much greater body of previously developed programs for use in on-line systems.

However, the system program and the apparatus- or data-manipulating programs must still be written in machine language or symbolic assembler language. In a large experimental group, some workers will be able to avoid learning the details of the computer. The lone research worker will be less lucky and must realize what he is getting into.

In the Industrial Context

Contrary to the usual situation in which industry applies tools developed in the laboratory, the successful use of computers to control complex processes was first achieved in the industrial environment. Even the earliest attempts at laboratory automation were made considerably after computer process control was well established in oil refineries and machinery manufacturing plants. Much of the early impetus toward development of sophisticated man-machine interfaces came from the requirements of the military and not from the laboratory. The laboratory has, nevertheless, profited from all these developments.

In the last few years the differences between industrial and laboratory "process control" have become clearer. First, laboratory automation is less expensive than its industrial counterpart because there need be so much less emphasis on reliability and redundancy of components. A day's downtime (time during which the machine is inactive) in a solid-state physics laboratory is much less costly than one in an automatic-crankshaft manufacturing facility.

Second, laboratory automation demands much greater flexibility. An industrial process, whether cake baking or paint spraying, is a reasonably stable enterprise. Changes of any significance are introduced on time scales of months and years. Not so with research, where procedures may be changed daily or weekly. This has put much greater demands on laboratory systems in terms of hardware and software versatility. Much of the demand for easily learned programming languages that can be used to control equipment may be traced to this need.

Some Applications

As should be expected with such a powerful technique, the applications have been varied and are distributed over most of the scientific disciplines.

A field which has profited considerably from the application of on-line computers is crystallography. A number of laboratories have incorporated computers into their neutron-diffraction and x-ray-diffraction apparatus. The step-by-step scanning procedures which characterize this work were natural candidates for automation. The techniques have been sufficiently successful so that the computer-driven diffraction equipment is now available commercially.

Research in neurophysiology has, similarly, been radically altered by the use of computers in many laboratories. Test subjects (animal and human) are stimulated by means of the computer; their responses are automatically recorded, correlated, and compiled into a growing body of data that is presented, in abstracted form, to the experimenter on a cathode-ray-tube display. The variety of experimental systems is astonishing.

In an allied area, research into the nature of perception and learning is being conducted, with the assistance of a computer, in a number of psychology laboratories. Again, the "tirelessness" of the computer in administering the programmed pattern and response sequence and simultaneously timing, measuring, and guiding has wrought a revolution in experimental practice.

In low-energy physics much special-purpose equipment (for example, pulse-height analyzers) has been replaced by the general-purpose computer. On-line systems are taking, sorting, and displaying data and making "live" comparisons of the data with theoretical expectations.

It may be fairly said that high-energy physics research, today, is totally dependent on the computer. Direct-reading spark-chamber and counter experiments are run "on-line," prodigious feats of data reduction being accomplished while the experiment is still in progress. Bubble-chamber films are scanned by automatic and manual apparatus, the data going directly to a computer. The computer, in turn, rapidly reconstructs the "event" being measured in order to provide the manual operators or the scanning instrument with timely information.

The computer has invaded the hospital. Medical diagnosis, monitoring of patients, and bioanalysis are all being done by computers at one place or another in the United States. The feasibility of pattern recognition by computer

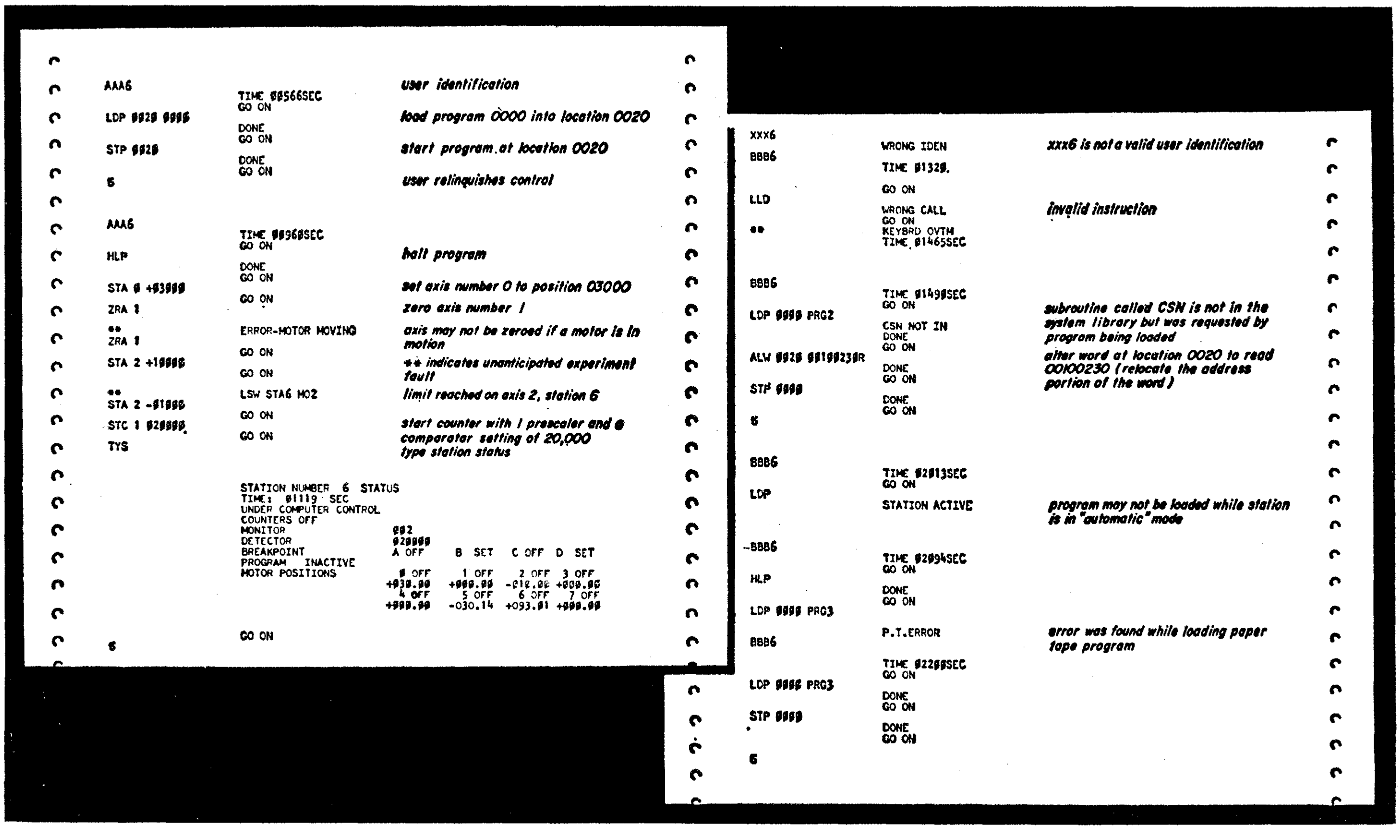

Fig. 4. Representative dialog at a control typewriter. The experimenter's shorthand commands appear at left. The Multiple Spectrometer Control System response to each command appears one line down in the next column. Italicized explanations are at right.

for cancer-cell detection is under active study. Chromosomes are recognized and paired; tumors are pinpointed; reminders of medications due are typed for nurses.

The number of applications snowball, being limited only by ingenuity and the availability of funds.

Multiple Spectrometer Control System

In order to particularize the concepts I have been discussing, I next describe a computer system that was first put into operation at Brookhaven National Laboratory in 1965 (*2*). It is called the Multiple Spectrometer Control System, and its function is that of simultaneously operating, controlling, and monitoring eight neutron spectrometers and one x-ray spectrometer at the laboratory's High Flux Beam Reactor. In addition to exercising direct digital control of all the mechanical and measurement processes, the computer performs all the complex data reduction computations necessary to effectively guide the separate experiments.

The nine spectrometers are used by independent research groups and must be individually controllable. The typical spectrometer has five motor-driven axes, two detectors, and various limit switches, error switches, ovens, Dewars, and so on which must be monitored. The programmed master control system (the "operating system") provides manual, semiautomatic, and completely automatic control for each unit. Each spectrometer can be operated from a local control station, from a typewriter (located in the computer room), or from the user's stored program.

It is the function of the control station to provide the direct interactions already familiar to users of manual controllers. The experimenter is provided with axis-position controls, scaler controls, decimal readouts and displays, and various general-purpose switches and signals. All control stations are identical and interchangeable.

The computer is located in a room that is central and accessible to the individual experimenters. It is from this room that all the semiautomatic and fully automatic operations are handled. Data recording and display for all the experiments are centered here. In addition, all system messages, error signals, and status reports are typed on one of the control typewriters for the operator's attention. The system supervisory control is so arranged that all the operating commands and station responses form an unambiguous, clocked log which proves quite useful in tracing experimental faults and in reconstructing test sequences about which there is uncertainty.

While the user has a great deal of flexibility in manipulating his spectrometer from the local station, some of the more complex setup and exploratory procedures can become rather tedious and time-consuming. When he wants to, he can take advantage of the broader range of facilities provided in the Multiple Spectrometer Control System by throwing the switch on the panel which puts his spectrometer under system control. With the spectrometer under the supervision of the control system, the user's method of control changes from manipulation of the specific functional switches to use of one of the control typewriters.

In its most elementary use, the typewriter control duplicates all the functions of the panel switches. It does, however, provide a more rapid and precise control of the axes and counters. In addition, it provides facilities for making various apparatus tests, status reports, and time checks.

In its broader use, the typewriter control allows the user to load into the computer memory a detailed operating program. This program may be the one which will run the experiment for the next several hours or days or, alternatively, may represent a series of extremely complex test manipulations which may then be selected, started, or stopped by means of the station's communication switches. The experimenter is thus able to select, at will, the degree of automatic control he wishes to employ at any stage of his experiment.

The user is free to organize his program in any way he chooses. Normally, however, he reserves a portion of his allotted space for storing data. When he has accumulated a block of information he calls on the Multiple Spectrometer Control System to write it out on magnetic tape, properly arranged and identified. The magnetic tape records, which are centrally accumulated for all users, are in a form directly suited for entry into the laboratory's central computer complex. Alternatively, the user may wish to reread a portion of the data previously stored. In particular, the system provides the facility which enables him to retrieve information and display it on the console cathode ray tube for immediate analysis. A light pen is also provided.

Beyond its normal role of providing the medium for the "dialog" between the user and his spectrometer experiment, the typewriter also keeps the system operator alerted as to the status of the whole experimental complex. The operator is automatically informed of such things as overtime usage, error conditions, and equipment malfunctions. The typewritten sheets, therefore, represent a complete and accurate record of total-system usage and performance. A representative section of an output page is shown in Fig. 4. (The italicized comments do not, of course, appear on the record. The user's "shorthand" entries appear at the left, and the system's response appears in the second column.)

The Human Factor

At the beginning of this article I spoke of the young graduate student's training in computer science. A recent survey (*3*) indicated that, among a representative group of scientists, more than 90 percent wanted their students to learn to program a computer, although only about half of the scientists themselves knew how to do so. The older the scientist was, the less likely he was to be skilled in these new techniques. (Only one in ten of the men over 50 was able to program a computer.)

One effect of this is that those least likely to understand these new techniques are the ones most likely to have to pass judgment on them. The instability of this situation has produced a corresponding polarization in views, the "pro's" being unreasonably pro and the "con's" being unreasonably con. Unfortunately, because the techniques of laboratory automation are radical, it is difficult to write about them and to adequately describe their effects on research. (One of my colleagues—now an enthusiast—recently confessed that when he first read about on-line computers he thought the idea was "just plain crazy.")

I encourage the reader to go see for himself.

References and Notes

1. The work discussed was supported by the U.S. Atomic Energy Commission. Rather than pepper this article with references I have chosen to list only a few tutorial and survey papers that will further serve to introduce this field to the reader and will, in turn, lead him to the rather extensive literature. These papers are as follows: R. J. Spinrad, *Ann. Rev. Nucl. Sci.* **14**, 239 (1964); D. A. Cooper, *Intern. Sci. Technol.* **1964**, No. 36, 20 (1964); R. J. Spinrad, *Progr. Nucl. Techniques Instrumentation* **1**, 221 (1965); S. J. Lindenbaum, *Ann. Rev. Nucl. Sci.* **16**, 619 (1966); J. F. Davis, *Intern. Sci. Technol.* **1966**, No. 60, 40 (1966); J. A. Jones, *IEEE (Inst. Elec. Electron. Engrs.) Trans. Nucl. Sci.* **14**, 576 (1967); "Conference on Use of Computers in Analysis of Experimental Data and Control of Nuclear Facilities, Argonne, 1966," *Conf. No. 660527, U.S. At. Energy Comm.* (1967); R. J. Spinrad, *Sci. Res.* **2**, No. 8, 38 (1967).
2. D. R. Beaucage M. A. Kelley, D. Ophir, S. Rankowitz, R. J. Spinrad, R. Van Norton, *Nucl. Instr. Methods* **40**, 26 (1966).
3. R. J. Spinrad, *Phys. Today* **18**, No. 12, 47 (1965).

Precision Measurement of the Acceleration of Gravity

Measurements of *g* have always made maximum use of the available technology in measurement of length and time.

James E. Faller

The measurement of the acceleration of gravity (*g*) has long been a matter of scientific interest. Almost from the beginning of time, men must have noticed that things move faster and faster as they fall. Over 300 years ago Galileo, in studying how things fall, discovered that the motion is one of constant acceleration. He showed that the distance a falling body travels from rest varies as the square of the time. To show this, he used an inclined plane to slow down the motion (to "dilute" gravity); by thus extending the time scale of his experiments, he was able to make quantitative measurements with the limited experimental means at his disposal. In his *Dialogues Concerning Two New Sciences* (*1*, p. 178), Galileo depicts for us the technology of his day as he describes his experiments concerning motion on an inclined plane:

> A piece of wooden moulding or scantling, about 12 cubits long, half a cubit wide, and three finger-breadths thick, was taken; on its edge was cut a channel a little more than one finger in breadth; having made this groove very straight, smooth, and polished, and having lined it with parchment, also as smooth and polished as possible, we rolled along it a hard, smooth, and very round bronze ball. Having placed this board in a sloping

The author is assistant professor of physics at Wesleyan University, Middletown, Connecticut.