CONTROL THEORY

In technology the most advanced applications of mathematics are in the design of machines that control themselves. The same methods are relevant to the control mechanisms of living organisms and societies

by Richard Bellman

Control theory, like many other broad theories, is more a state of mind than any specific amalgam of mathematical, scientific or technological methods. The term can be defined to include any rational approach used by men to overcome the perversities of either their natural or their technological environment. The broad objective of a control theory is to make a system—any kind of system—operate in a more desirable way: to make it more reliable, more convenient or more economical. If the system is a biological one, the goal may be to understand how the system works and to reduce pain and distress.

In this article I shall mainly discuss control theories that have some explicit mathematical content, but it is clear that some of the most interesting control problems arise in fields such as economics, biology and psychology, where understanding is still notably limited. To prove that he understands, the scientist must be able to predict, and to predict he requires quantitative measurements. To make predictions that are merely qualitative, such as the prediction that an earthquake, a hurricane or a depression will occur in the near future, is not nearly so satisfying as a similar prediction associated with a specific time and place.

The ability to make a quantitative prediction is normally a prerequisite for the development of a control theory. In order to make quantitative predictions one must have a mechanism for producing numbers, and this requires a mathematical model. It might seem that the more realistic the mathematical model is, the more accurate the prediction will be and the more effective the control. Unfortunately, diminishing returns set in rapidly. The real world is so rich in detail that if one attempts to make a mathematical model too realistic, one is soon engulfed by complicated equations containing unknown quantities and unknown functions. The determination of these functions leads to even more complicated equations with even more quantities and functions—a tale without end.

The richness of the problems presented by modern civilization have led to the study of control theory on a broad front and to the development of a large variety of control systems. Although this development began long before the invention of electronic computers, the explosive growth of control theory dates from the appearance of these devices soon after World War II. For the past 20 years control theory and computers have grown side by side in an almost symbiotic relation. Without the computer most of the advanced control systems used in the military domain, in space technology and in many branches of industry could not have been developed, and without the computer they certainly could not be effectuated.

In industry, control theory, implemented by the computer, is now widely used to regulate inventories, to schedule production lines and to improve the performance of power stations, steel mills, oil refineries and chemical plants. It is estimated that about 500 computers specially designed for process control are now installed or on order. Five years ago scarcely a dozen such machines were in service.

A majority of process-control computers still operate in "open loop" fashion, which means that they monitor the process variables, analyze them in a search for possible improvements and present their recommendations to a human operator for action. In a growing number of plants, however, the loop has been closed. The decisions of the computer are directly linked to the process controls so that adjustments are made automatically.

In Rotherham, England, to select one remarkable example, a new steel plant built by Tube Investments Ltd. has been provided with three large digital computers arranged in a chain of command. The computer at the top is "off line" and is used for production planning. It receives customers' orders and classifies them according to the composition of the steel and the form of the finished product. It then calculates an efficient three-week program for the steel furnaces and rolling mills and keeps track of the program's progress. The second computer takes over when a billet has been produced by one of the furnaces and produces a full set of instructions for its subsequent treatment. This computer is "on line" and actually supervises the rolling mill. The third computer has the single task of receiving measurements of billet size and computing how they should be cut to minimize waste. This kind of integrated operation from receipt of the customer's order to final billing has become the goal of many manufacturing firms.

The industrial, military and space-flight control problems that have been presented to mathematicians for solution in the past 25 years have brought about

CLOSED-LOOP CONTROL is implemented by a computer (*console in foreground in photograph on opposite page*) in the operation of a 650-million-watt, coal-fired power plant at the Paradise Station of the Tennessee Valley Authority near Drakesboro, Ky. One of the largest power plants in the world, the station has two such computer-controlled units. In a closed-loop system the computer directly adjusts the process variables.

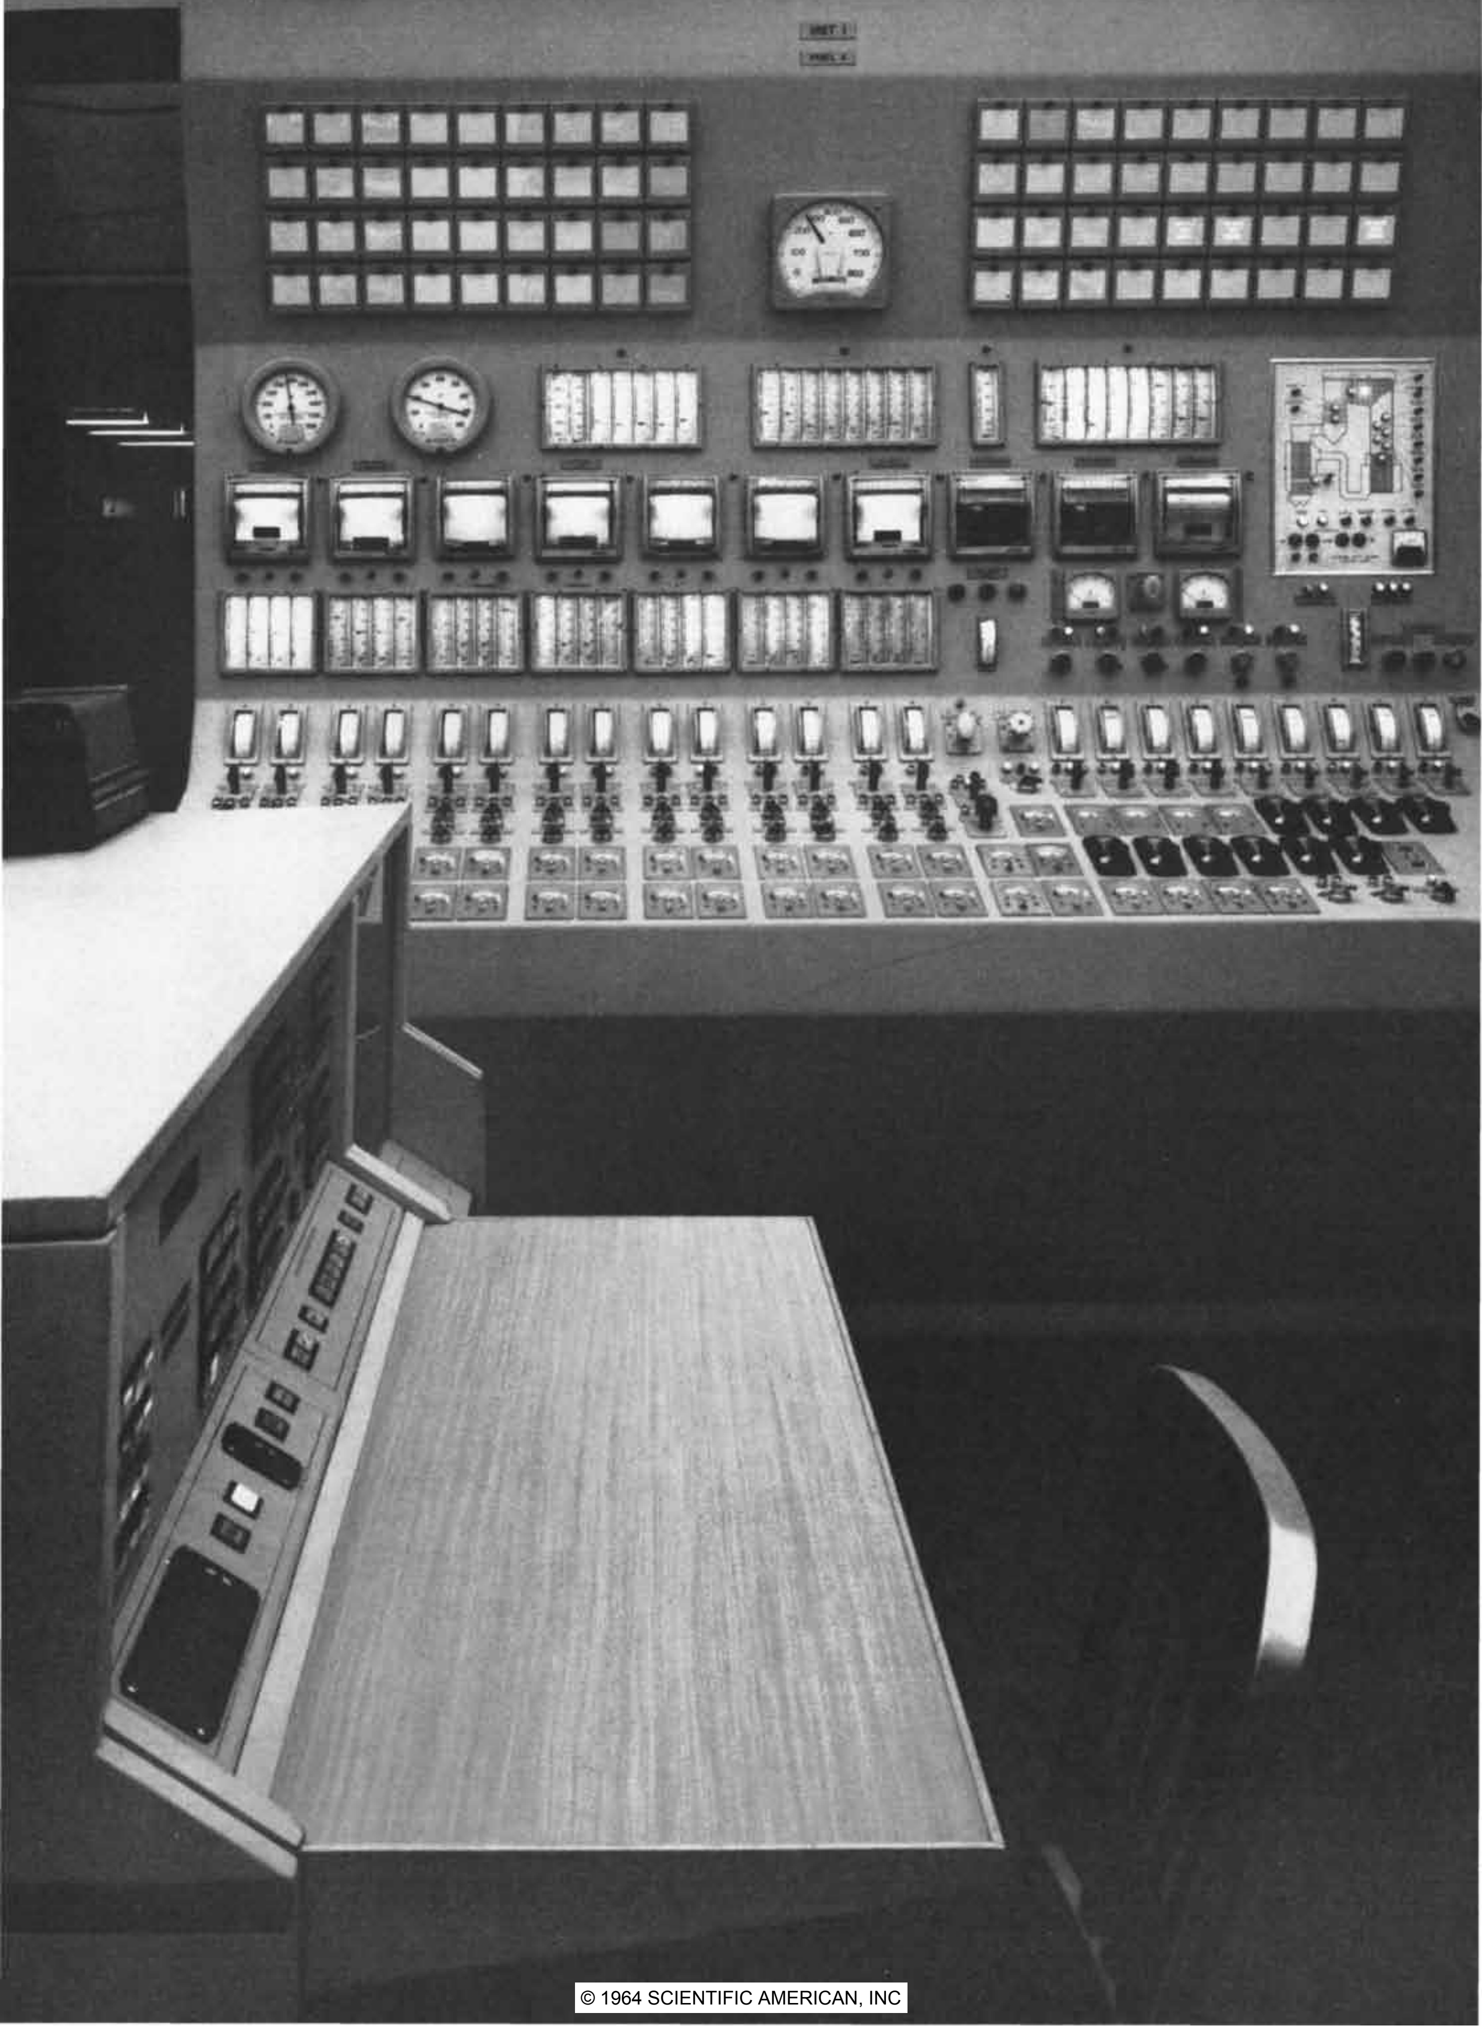

the revitalization of a number of moribund mathematical disciplines and have led to the creation of new theories of considerable intrinsic interest. Since control and stability are intimately related, mathematical theories devised in the 18th and 19th centuries to study such matters as the stability of the solar system have been dusted off, refurbished and applied to many problems of more current interest. These theories have included highly abstruse conceptions of the great French mathematician Henri Poincaré and the Russian mathematician A. M. Liapunov, which are now routinely employed in control studies.

The most challenging control problems encountered in science, technology, economics, medicine and even politics can be described as multistage decision processes. Traditionally they have been treated on the basis of experience, by rule of thumb and by prayerful guesses. The basic task is to determine feasible and reasonable courses of action based on partial understanding and partial information. As more information is obtained one can expect to do better, but the crux of the problem is to do something sensible *now*.

A familiar problem characterized by partial understanding is that of maintaining a healthy national economy—of avoiding a depression on the one hand and inflation on the other. A variety of regulatory devices are available for achieving the desired control. One device is to regulate the interest rate on loans. If inflationary trends develop, the interest rate is raised and money gets tighter; if a depression impends, interest rates are decreased, the investment rate rises in response and more money enters circulation.

The policy that should be pursued depends critically on what is occurring in the system at the moment. For one to know what is going on requires a feedback of information. The concept of feedback control is now familiar to almost everyone. It means an automatic regulating linkage between some variable and the force producing it. One of the earliest applications of feedback control in technology is the governor used by James Watt on his steam engine. Even earlier Christian Huygens had devised what might be called a static feedback system to regulate the period of a clock pendulum [*see top illustration on next page*].

Usually it is a combination of complexity and ignorance (in polite circles referred to as "uncertainty") that forces one to employ feedback control. If, for example, the workings of the economy were as fully understood, let us say, as the movements of the planets, one could predict well in advance the behavior of producers and consumers; one could predict such things as the effects of population growth and the consequences of introducing new goods and

CHEMICAL PROCESS CONTROL SYSTEM in which a digital computer (*foreground*) exercises closed-loop control is shown in the Bishop, Tex., plant of the Celanese Chemical Company. The plant converts petroleum gases to acetic acid, acetaldehyde and other chemicals that are used in paints, plastics, fibers, drugs, cosmetics, fuels and lubricants. This computer and those at Paradise Station of TVA were built by Bunker-Ramo Corporation. Closed-loop computer systems are also installed in steel-rolling mills.

services. On the basis of this knowledge one could compute and announce the desired interest rate for years ahead. It should be observed that one would then have to reckon with the consequences of publishing the rates in advance, because producers and consumers would include the *future* rates in their *current* economic decisions.

In actuality one must adopt a wait-and-see policy. One observes the economic scene for a period of time and draws conclusions about current trends. On the basis of these conclusions the interest rate, or some other control lever, is changed. One hopes that the action is well timed, or in phase. The matter of the timing of external influences is of crucial importance in control theory. Anyone who has pushed a swing is familiar with the consequences of applying the impulse a fraction of a second too soon or too late.

Since complexity and uncertainty abound in modern control problems, the use of feedback has become routine. In fact, it is sometimes forgotten that control problems are still solved without direct, or active, feedback. This is the case, for example, when an automatic machine tool is set up to turn out a number of identical parts. It is assumed that the control problem is completely solved in advance. In practice, of course, the parts vary slightly and finally exceed the prescribed tolerance, whereupon the machine is readjusted. This is feedback control after the fact.

In the newest machine tools the dimensions of the workpiece are monitored continuously, and feedback control is employed to regulate precisely the amount of metal removed. In this way it is possible to turn out parts that are as nearly identical as one might wish. For such jobs digital computers can be used, but they are not essential.

The computer is essential when complex decisions must be made at high speed, as in the launching of a space vehicle. This is a multistage decision process whose solution is contingent on information acquired and fed back to the control system as the process unfolds. A computer, either aboard the rocket or on the ground, is essential for making a succession of decisions as rapidly as may be necessary. Such a computer is said to be operating in "real" time, because it must keep pace with the problem being solved.

A process-control computer installed in a refinery must also operate in real time, but the time available for making a decision may be 10 or 100 times longer than that available for rocket

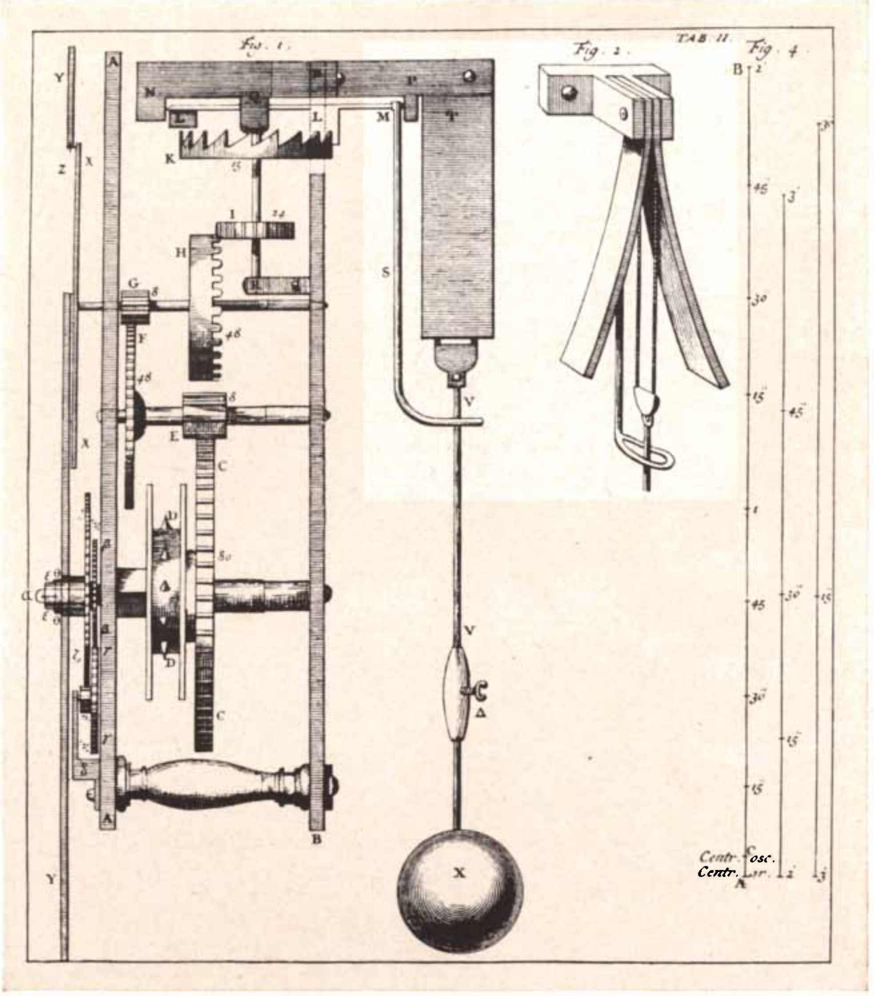

FEEDBACK CONTROL for a clock pendulum was invented in 1673 by the Dutch mathematician Christian Huygens. The curved metal strips on each side of the pendulum cords (*seen in perspective at right and labeled "T" in the side elevation of the clockwork at left*) were designed to make the period of the pendulum constant regardless of the length of its arc. The rod *S* moved with the pendulum, transmitting its motion to the clock.

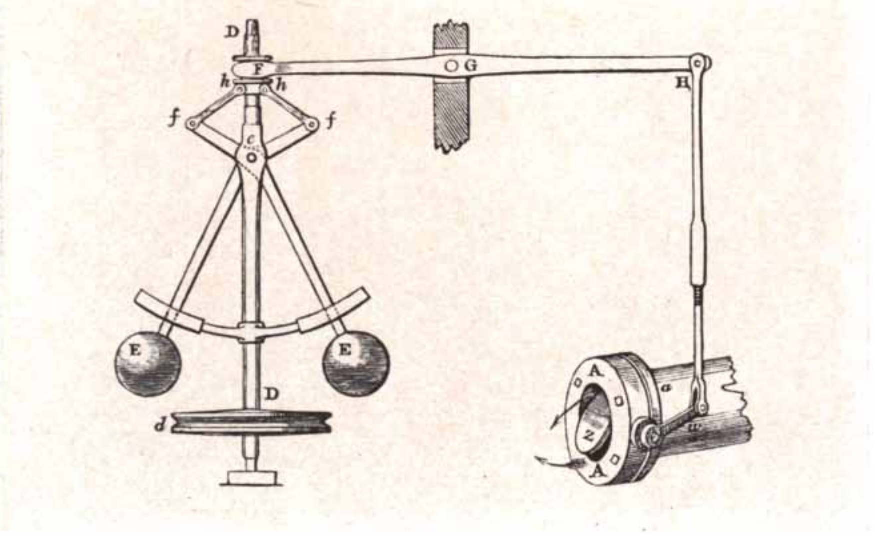

FLYBALL GOVERNOR (*left*), one of the earliest automatic control devices, was invented by James Watt to govern the speed of the steam engine. As the engine speeds up, the rod (*D*) on which the balls (*E*) are mounted spins faster, causing the balls to fly outward. This in turn closes the butterfly valve (*Z*), decreasing the supply of steam to the engine (*not shown*) and slowing it. A fraction of the output of the engine goes into the rotation of the flyballs; a fraction of this fraction is fed back to govern the speed of the engine.

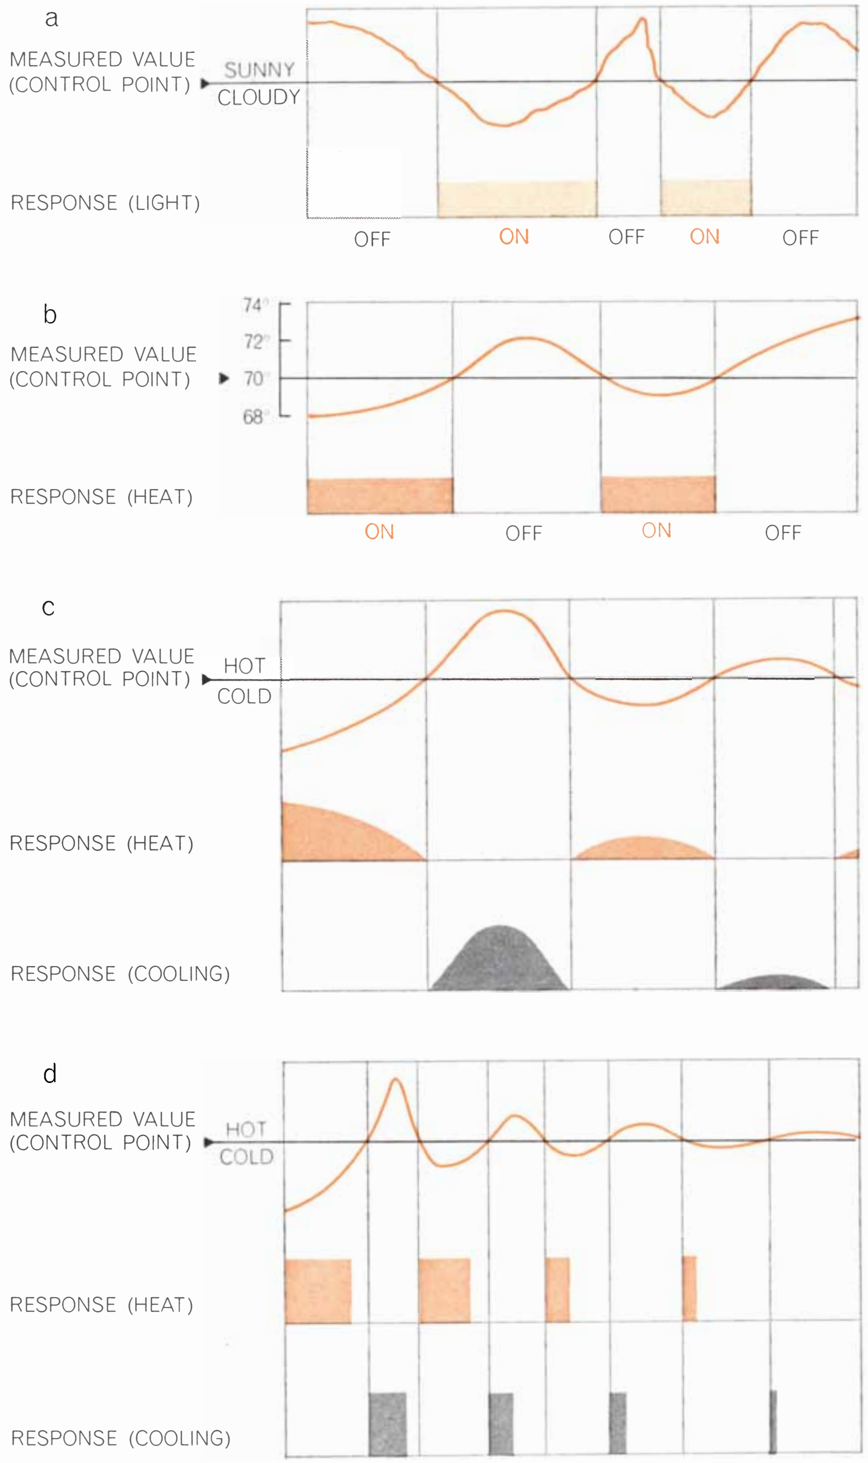

FOUR CONTROL SYSTEMS show how a measured variable can be brought under increasing refinement of control. Diagram *a* depicts a simple on-off response to a measured value, such as turning on the lights in a room when the sun goes behind a cloud. The measured value is not regulated and feedback is not employed. In *b*, which represents a typical home-heating system, on-off response is combined with feedback. When the temperature falls below the desired value, the furnace goes on, but since no cooling system is provided the room temperature may climb higher than desired on a sunny afternoon. The system in *c*, which could represent the heating of a chemical reaction vessel, provides both heating and cooling. The response is graduated so that as the control point is approached the rate of heating or cooling is reduced. The control problem in *d* is the same as in *c* but two modifications have been introduced to improve the speed and accuracy of control. Heating and cooling are not graduated but operate at a constant high rate when called for. This is known as "bang-bang" response. In addition a computer in the control system measures the rate of change in the controlled variable, takes account of the time lag in the temperature-recording mechanism and shuts off the heating or cooling before the control point is reached. Thus oscillation, or "hunting," of the system is damped out quickly.

guidance. On the other hand, a process-control computer may have to deal with 10 or 100 more variables than the rocket computer. And it may have to review a lengthier sequence of logical alternatives before making a decision.

What tools does the mathematician have for trying to control a multistage decision process? The conventional approach can be labeled "enumerative." Each decision can be regarded as a choice among a certain number of variables that determine the state of the process in the next stage; each sequence of choices defines a larger set of variables. By lumping all these choices together the mathematician can "reduce" the problem to a Newtonian one of determining the maximum of a given function.

It would seem simple enough to maximize a reasonably well-behaved function; using the familiar technique of calculus, the mathematician takes partial derivatives and solves the resulting system of equations for the maximizing point. Unfortunately the effective analytic or numerical solution of many equations, even apparently uncomplicated linear ones, is a difficult matter. By itself this is nothing more than the "curse of dimensionality," with which physicists have had to live for many years; significant results can be obtained in spite of it.

There are, however, more serious difficulties. In many cases the solution is a boundary point of the region of variation. This corresponds to the constraints, or restrictions, of real physical and engineering systems. When this is so, calculus is often inadequate for discovering maximum and minimum points and must be supplemented by tedious (and usually impossible) hunt-and-search techniques. Finally there is the frequent complexity that the outcome of a decision is not explicitly determined but is itself a random variable. The process is then said to be stochastic. Here to an even greater extent any simple enumerative technique is doomed by the vast proliferation of possible outcomes at every stage in the process. One cannot enumerate "all" possibilities and choose the best—not when there are 10^{50} or 10^{100} possibilities.

Has the mathematician now reached the end of his resources? Not if he will step back and ask himself if he has understood the nature of the solution he is seeking. How, he must ask himself, is the form of the solution influenced by the physical properties of the system? In other words, the mathematician cannot consider his problem solved until

$$\frac{\sqrt{\textbf{if}}}{\textbf{maybe}}$$

Solve for profit.

(Fischer & Porter can. And does. In processes all over the world. Yours too. If you'll ask.)

Processing is, at best, a matter of working with the best approximations you can get.

To lick process problems, you need three things: good instrumentation; plenty of experience, preferably world-wide; and the ability to use that instrumentation and experience to build good process control systems.

Take instrumentation, for instance, and flowmeters in particular. (If any one process variable can be considered as "basic," flow is the one.) Fischer & Porter builds more flowmeters than anyone else in the world. And more kinds of flowmeters. And the world's largest flowmeters. And the most accurate flowmeters. This, obviously, puts us one up on your process problems.

Or consider facilities and experience around the world. F&P has plants in nine countries overseas, engineering offices in 37 others. And these people have obviously been working at their trade, since over 25% of our problem-solving is done overseas. All this experience is automatically yours when you call in Fischer & Porter.

And control-systems experience? Just last year, Fischer & Porter built the world's largest petroleum blending system. And the controls for the world's largest filtration plant. (And they weren't even the biggest orders in the shop.)

Q. E. D. When you need a process control system, Fischer & Porter can help. Call us.

FISCHER & PORTER COMPANY

A world-wide control instrument company with manufacturing plants in Australia, Canada, England, France, West Germany, Italy, Mexico, the Netherlands, Spain and the United States.

County Line Road • Warminster, Pennsylvania • 215-675-6000 he has understood the structure of the optimal policy. Let me explain.

We have seen that in the conventional approach the entire multistage decision process is regarded as if it were a one-stage process. Thus if the process has N stages and there are M decisions to be made at each stage, the conventional approach envisages a single-stage process in MN dimensions. What one would like is to avoid this multiplication of dimensions, which stifles analysis, fogs the imagination and inevitably impedes computation.

The alternative approach—the policy approach—places emphasis on the characteristics of the system that determine the decision to be made at *any* stage of the process. In other words, instead of determining the optimal sequence of decisions from some *fixed* state of the system we wish to determine the optimal decision to be made at *any* state of the system. Only if we know the latter state do we understand the intrinsic structure of the solution.

The mathematical virtue of this approach lies first of all in the fact that it reduces the dimension of the problem to its proper level, which is the dimension of the decision that confronts one at any given stage. This makes the problem analytically more tractable and computationally much simpler. In addition this approach provides a type of approximation ("approximation in policy space") that has a unique mathematical property ("monotocity of convergence"). This means that each successive approximation improves performance [*see illustration on page 194*].

The name I proposed some years ago for this policy approach to multistage decision processes is dynamic programming. One of its goals is the determination of optimal feedback control. The adjective "dynamic" indicates that time plays a significant role in the process and that the order of operations may be crucial. The approach is equally applicable, however, to static processes by the simple expedient of reinterpreting them as dynamic processes in which time is artificially introduced.

Dynamic programming has given rise, in turn, to subsidiary and auxiliary control theories that go by a variety of names: theories of stochastic and adaptive variational processes, theories of Markovian decision processes, theories of quasi-linearization and invariant embedding. They cannot be explained in a few words; I mention them merely to indicate how control theory has branched and developed in recent years.

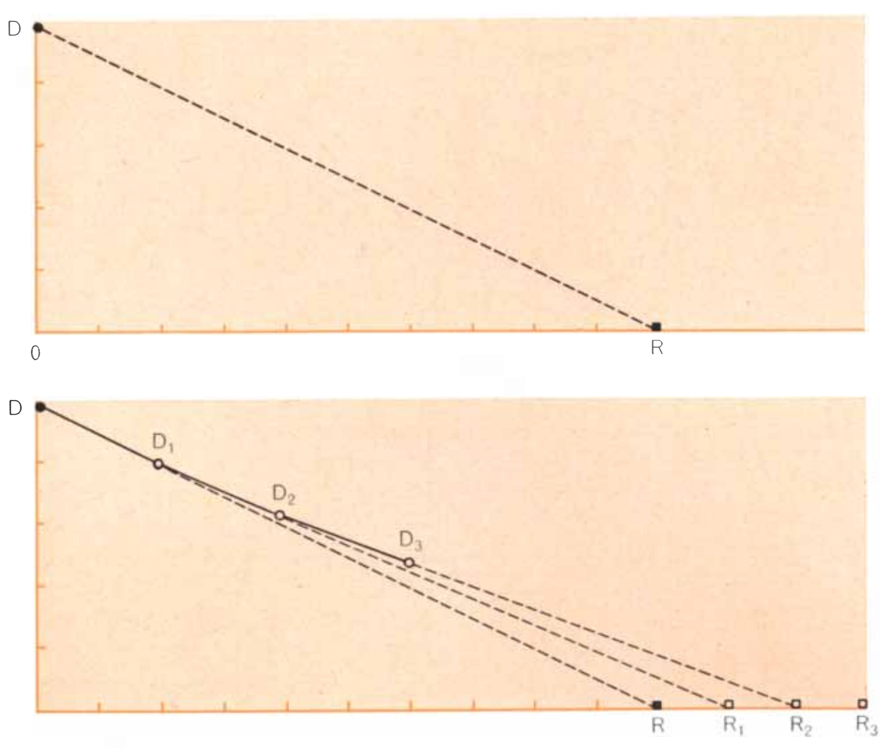

PURSUIT PROBLEM can be solved by adoption of a simple policy that lends itself to computer implementation. The problem is to find the path traced by a dog (D) chasing a rabbit (R). At the outset (*top*) the rabbit is 100 feet from 0 and the dog is 50 feet from 0. The dog runs at 22 feet per second, the rabbit at 11 feet per second. The dog always continues in a particular direction for one second. After the first second the dog has reached D_1 and the rabbit R_1 (*bottom*). To determine the point D_1 a straight line is drawn between D and R and 22 units are measured along it. Similarly, D_2 is determined by connecting D_1 and R_1, and so on. The resulting path approximates the one taken by the dog and can be refined by changing the direction of the path at shorter and shorter intervals.

Let me now illustrate how the adoption of a policy can simplify a problem that otherwise would be hard to handle on a computer. (Complex versions of the problem cannot be handled without a computer.) Consider the problem of a hotel manager who wants to provide chairs for a group of people in a room. He has a helper who carries chairs with ease but who cannot count. What does the manager do?

He employs the primitive and powerful concept of equivalence, together with feedback control. He tells the helper to keep bringing chairs until everyone in the room has a chair. This sequential procedure guarantees that each person will have a chair, without ever determining how many people or chairs there are. Furthermore, if some chairs are defective, a simple modification guarantees that everyone will eventually have a sound chair.

Consider next the case of an elderly woman whose memory is failing. It irritates her to have to hunt through her wardrobe for various items of clothing when she dresses in the morning. She could create a filing system, complete with a written index, but this would be a lot of trouble. Instead she solves her problem by putting a complete outfit in every available drawer.

In both cases the solution is quite "simple," but it is not necessarily obvious. Both ideas are currently used in programming computers to solve complex problems. The first is used in certain simulation processes and in Monte Carlo calculations. The second is used for retrieving key items of information from a very large computer memory. Since the items are needed frequently they are stored in several places, thus considerably reducing retrieval time.

I might add that many mathematicians have the nagging suspicion that the universe is much simpler than it appears in their complex mathematical models. It is not easy, however, to capture the fresh view required for the simple approach. In the use of computers, changes in viewpoint such as the two just mentioned have time and again changed an impossible problem into a possible one, a merely difficult one into a routine exercise.

The next example is chosen to show how the concept of policy can not only simplify a multistage decision problem

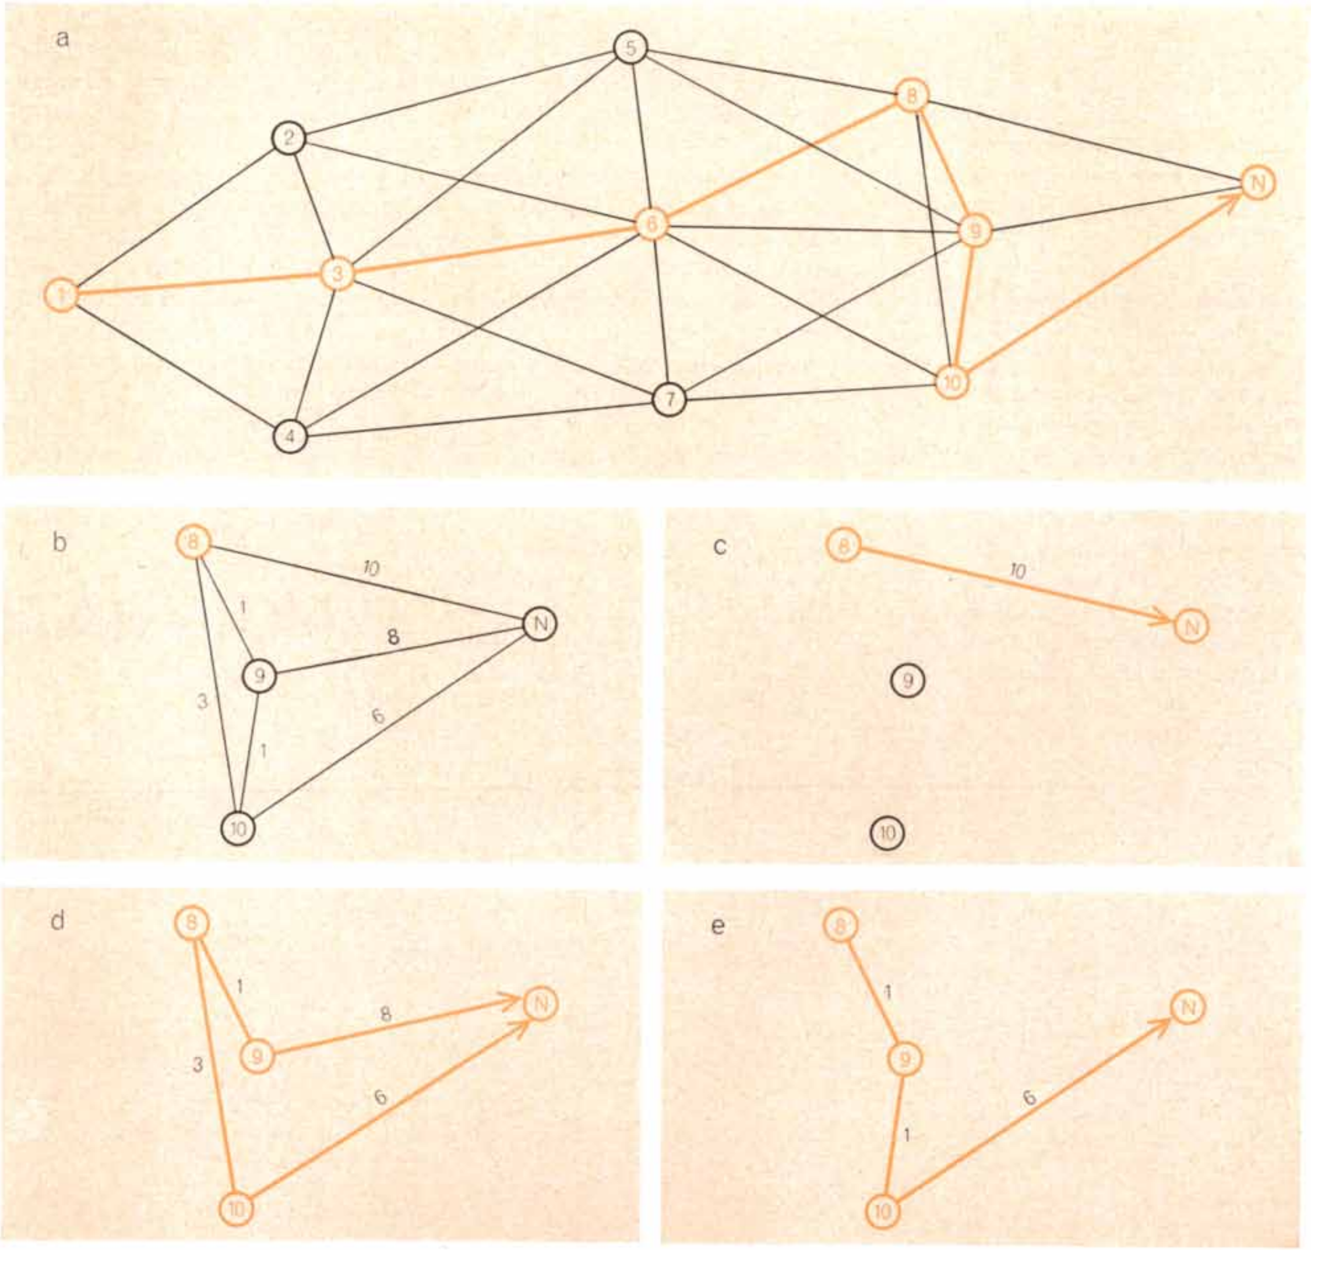

ROUTING PROBLEM is commonly met in control theory and was generally difficult to solve before computers and special programming methods were developed. The problem shown here is to find the path from *1* to *N* that requires the minimum time. The circles represent towns that are connected by a network of roads and the travel time between all pairs of towns is known. The traditional approach to such a problem was simply to enumerate all possibilities. This quickly leads to a "race against *N*." In the example as drawn, in which *N* is only 11, the number of different routes from *1* to *N* (with no backtracking) is more than 10,000. If *N* were 30, a high-speed computer would need more than 100 hours just to enumerate all the possible routes. One way to make such problems tractable is to use "dynamic programming," which depends on a selection of "policies." The virtue of such policies is that they can be applied from any point *i* in the network and thus satisfy the injunction: Do the best you can from where you are. The four smaller diagrams show how policies are selected assuming that *i* is point *8*. The travel times, in hours, for the various routes from *8* to *N* are shown in *b*. The initial policy (*c*) is to go directly from *8* to *N*, which takes 10 hours. The second policy (*d*) is to make one stop, which provides two alternative routes of nine hours each. The third policy (*e*) is to make two stops, which provides the minimum-time route. If point *1* were selected as the initial point *i*, the same procedure would be followed, but since this particular network does not provide a direct path from *1* to *N*, the first policy that could be examined would be one with the least number of stops, in this case three. There are, of course, ways of formulating this policy approach in terms of a computer program. The equation for solving the problem is

$$f_i = \min_{j \neq i} [t_{ij} + f_j],$$

in which f_i is the minimum time from any point *i* to *N*; t_{ij} is the time required to go from *i* to any other point *j*, which may be *N* itself, and f_j is the minimum time to go from *j* to *N*. This dynamic programming equation is solved by successive approximations—"approximations in policy space"—in which each successive approximation improves the result. The equation can be solved numerically for networks of several hundred points by hand in a few hours and by electronic computer in a few seconds. The equation determines both the optimal policy ("Where does one go next?") and the minimum time. Moreover, the equation embodies all the mathematical power of the classical calculus of variations.

but also yield numerical answers. To follow this example the reader must refer to the illustration on page 192, which shows the path of a dog chasing a rabbit. The dog is initially at D; the rabbit is at R and is running to the right along the x axis. If the dog always heads straight toward the rabbit, what curve does the dog follow?

This is a standard problem in the theory of differential equations, but the nonmathematical reader would hardly be edified if he were given the explicit form of the solution. To understand it requires a certain amount of mathematical training. This is strange when one thinks about it; the dog solves the problem without hesitation, although of course he does not get numerical answers.

It is easy to obtain a good approximation of the shape of the curve in the following way. Let us assume that the dog can run 15 miles per hour, or 22 feet per second, and that the rabbit can run just half as fast, or 11 feet per second. The rabbit is originally 100 feet from 0 and the dog is 50 feet from 0 at a point perpendicular to the x axis. Assume now that the dog continues in any particular direction for one second at a time. At the end of one second the dog has reached D_1 and the rabbit R_1. Another second later the dog has reached D_2, the rabbit R_2 and so on.

The point D_1 is determined by connecting D and R with a ruler and measuring 22 units along it. Similarly, D_2 is determined by connecting D_1 and R_1 and repeating the same measurement. The process is continued until the distance between dog and rabbit is closed. (We will ignore the fact that the closing stages are made a little messy by this method.) The broken-line path is a simple approximation of the actual path traversed by the dog. It is evident that the approximation can be made as close as desired by carrying out the change in the dog's direction at shorter and shorter time intervals, say every hundredth or every millionth of a second. By hand computation this would be increasingly tedious, but an electronic computer can do the job easily in a matter of seconds.

More sophisticated versions of this problem occur in the determination of optimal trajectories for space vehicles. In some of these cases the "rabbit" is imaginary and the problem is to determine where to point to achieve a desired course; in other cases the "rabbit" is real enough—another craft or a planet, perhaps—and establishing exactly where it is provides a further significant complication.

The point I wish to emphasize is that one can obtain a solution to the original problem by concentrating on the original process. One merely follows the instructions for what to do at every point in time and space. In mathematical terms, one carries out a policy.

The importance of this from the standpoint of control theory is manifold. In the first place, it is easy to use computers to implement policies. In the second place, the mathematical level is more fundamental, deeper and yet relatively uncomplicated by symbol manipulation. Policies are invariably simpler than time histories. Much more emphasis is now placed on the formulation of the problem. The idea is to take full advantage of the structure of the process in order to describe it in a most convenient analytical fashion and in order to make clear the structure of the optimal policy. One tries to avoid any routine description in terms of complicated equations that do not easily yield to numerical approach. One does not try to fit every new type of decision process into the rigid mold of 18th-century mathematics. This is the policy concept behind dynamic programming.

With this concept, which recognizes the resources of the digital computer and accepts it as an ally, one can easily and quickly obtain the numerical solution of control problems in many different fields that defied even the most resourceful mathematicians 20 years ago. The new approach has made it possible to solve formidable problems in trajectory analysis, process control, equipment replacement and inspection procedures, communication theory, the allocation of water resources and hydroelectric power, the use of forest resources and investment planning—to mention only a few important areas.

Beyond this, the concept of policy can readily be applied to study the more difficult and realistic classes of decision processes involving uncertainty and learning. I have already referred to the former as stochastic processes; the latter are known as adaptive control processes.

The dog-rabbit pursuit process was an example of a deterministic process in which the basic mechanisms are fully understood and it is "merely" a matter of devising a suitable procedure for solving it. Thus we assumed that the positions of both the dog and the rabbit can be precisely observed and determined at each instant, that the

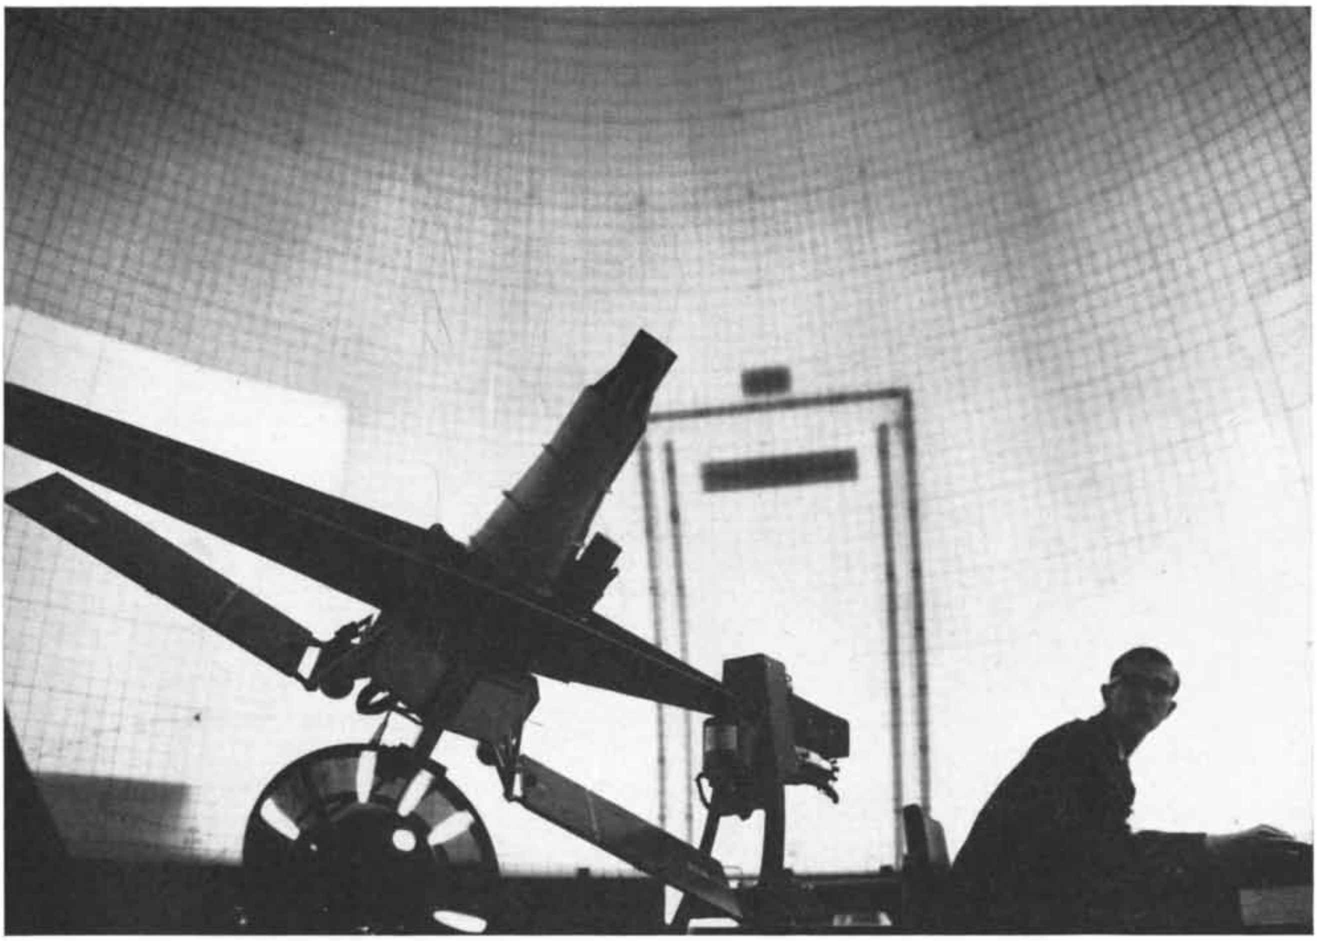

He's not a doctor. He's an essential.

The Jet Propulsion Laboratory has more than its share of Ph.D.s. But without the bachelors and masters, the doctors' hands would be tied.

Almost all engineering disciplines are involved in space exploration. Systems analysis, communications, guidance and control, propulsion, space mechanics, basic research . . . it's a long, long list. It has to be. The moon and planets are a long, long way away.

If you're an engineer, if you'd like to work in the wide open space, you may have the essential talent JPL is looking for. The way to find out is to write to us about you. Do it soon, why don't you?

Send your complete resume to:

JET PROPULSION LABORATORY
4802 Oak Grove Drive, Pasadena, California
Attention: Personnel Dept. 9

"An equal opportunity employer." Jet Propulsion Laboratory is operated by the California Institute of Technology for the National Aeronautics & Space Administration.

speeds remain constant, that nothing distracts the animals and so forth.

Even this idealized situation leads to difficult enough mathematics and a plethora of unsolved problems, as the study of classical celestial mechanics has demonstrated. After centuries of observation and computing, the positions of the planets next year or 10 years from now cannot be predicted with the accuracy desired. There will be a significant discrepancy between their predicted positions and those actually observed. A more upsetting problem for long-term planners is that no one knows if the solar system is completely stable.

It is obvious that if the idealized situation of perfect information and prediction cannot be found in the planetary motions, it can hardly be found in problems of trajectory optimization, satellite control and space rendezvous, much less in chemical process control, economic planning and medical diagnosis. In practice we are constantly using fallible devices for sensing and measuring, for processing, storing and retrieving information and for carrying out control decisions. Thus at every step we introduce error: error in observation, in calculation, in decision, in operation and even in the evaluation of outcomes.

The concept of a policy involving feedback control is ideally designed to handle the certain uncertainties of the actual world. By means of dynamic programming the injunction "Do the best you can in terms of where you are" (which is eminently sound common sense) can be readily translated into algorithms, or sets of rules, for the rigorous formulation and numerical solution of stochastic control problems.

When we turn to adaptive control processes we find a still higher order of uncertainty. In the stochastic case it is tacitly assumed that we know the detailed structure of the system we are studying, that we know various causes and various effects and, perhaps most essential of all, that we know what we want to do. In the case of an adaptive control process none of these assumptions may be valid.

Virtually all the unsolved major health problems can be regarded as adaptive control problems. Since no one knows the causes of cancer, coronary disease or mental illness, therapies aimed at control are necessarily based on a wide variety of hypotheses. This explains, of course, why so much caution must be exercised in treating patients. In the study of our national economy no one knows exactly what

WHY ARE 19 SYLVANIA LABORATORIES IMPORTANT TO AN ENGINEER WHO NEEDS ONLY ONE JOB?

These laboratories are linked together, within the integrated, nationwide complex of Sylvania Electronic Systems. Their activities traverse the entire gamut of advanced electronics. As a member of this variegated technological family, the individual engineer finds his advancement avenues and opportunities for interdisciplinary exposure have been multiplied, virtually by a factor of 19!

Cross-communications and cross-fertilization are excellent. Technical achievements, originating in, say, Mountain View, California—or in Williamsville, New York—or in Waltham or Needham, Massachusetts—are available to staff members across the Division. Examples of recent SES accomplishments range from survivable ground command/control networks for unmanned Minuteman sites (a Division-wide prime contract) ... to a laser breakthrough (result of a company-sponsored program). The laser wide-band modulator device developed has a theoretical capacity of handling simultaneously on one communications channel—more than 100,000 telephone conversations or 100 TV broadcasts.

Just how extensive SES' wealth of state-of-the-art knowledge is may be judged by the following:

In 1963, Sylvania ranked 19th among 500 DOD contractors (from all technical fields) in experimental, developmental, test and research work.

BEHIND THE SES ENGINEER ARE THE STABILITY AND FACILITIES OF GENERAL TELEPHONE AND ELECTRONICS.

The SES engineer has a special relationship with the parent corporation, General Telephone and Electronics, because SES acts in a Systems Management and Systems Program Development capacity for the entire corporation on major government contracts. In addition to operations which concentrate exclusively upon systems work in its most advanced military electronics aspects, SES headquarters houses the Division's Applied Research Laboratory which originates and advances electronic techniques required for tomorrow's more sophisticated systems.

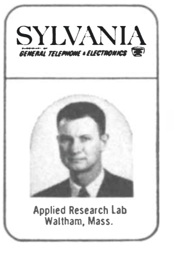

WOULD YOU WORRY ABOUT "TECHNOLOGICAL OBSOLESCENCE" UNDER THESE CIRCUMSTANCES? We think not. Certainly not at all, if you took advantage of SES' professional self-development training programs such as in-plant seminars, symposiums and conferences plus university post graduate studies and professional society meetings.

SYLVANIA ELECTRONIC SYSTEMS
Government Systems Management
for GENERAL TELEPHONE & ELECTRONICS

An Equal Opportunity Employer

will happen if taxes are cut or if military spending is reduced. Furthermore, there is a continuing controversy even about what constitutes a desirable economic condition.

When faced with an adaptive control problem, one expects to learn more about the system as time goes on and to modify one's policies accordingly. All major decision processes in life are adaptive control processes. It should not be surprising, therefore, that biological evolution has equipped animals to deal more or less successfully with adaptive control problems. Deterministic and, to a degree, stochastic control problems can be handled by animals on the basis of instinct. Instinct can be described as feedback control of a deterministic type. The same stimulus produces the same reaction, regardless of what else has changed in the environment.

To handle adaptive control problems the higher animals are equipped with something we identify as "intelligence." In fact, intelligence can be defined as the capacity to solve, in some degree, an adaptive control problem. Intelligence manifests itself by adaptation, by flexible policies. It is difficult, of course, to draw a sharp dividing line between instinct and intelligence. It is probably better, then, to call every type of feedback behavior "intelligence" and subsequently distinguish between levels of intelligence.

Norbert Wiener, the eminent mathematician who died last winter, formulated the provocative idea that it should be possible to develop a unified theory of feedback control applicable both to living organisms and to machines. To express this idea he coined the term "cybernetics." It was his hope, shared by others, that techniques used so successfully in control engineering could be applied to biomedical problems (for example the design of artificial human organs) and also that research into neurophysiology might provide valuable clues in the design and study of communication systems, computers and more general control systems of all kinds. But as mathematicians, physiologists and engineers explore the subtle difficulties of dealing with large-scale systems—living and nonliving—of different degrees of complexity, it seems less and less likely that any single "cybernetic" theory will serve all purposes. Nonetheless, for those who want to understand both modern science and modern society, there is no better place to start than control theory.