BYTE 1: June 1979

The Nature of Robots

PART 1: DEFINING BEHAVIOR

About the Author

William T Powers has been exploring the meaning of control theory for studies of human nature since 1953, when he was working as a health physicist at the University of Chicago. Since that time he has spent a number of years (to 1960) in medical physics, and then another 13 (to 1975) as Chief Systems Engineer for the Department of Astronomy at Northwestern University. His occupation has been designing electronic, optical, and mechanical systems for science. Powers' book, Behavior: The Control of Perception (Aldine, 1973) was quite well received. At the moment he consults in one-of-a-kind electronics.

In This BYTE

William T. Powers has a control theory approach to the simulation of human behavior. However, before we can simulate human behavior in a robot, we must determine what behavior is. William Powers takes a look at behavioral actions as he explores **The Nature of Robots**. Page 132

This article appeared in BYTE magazine, volume 4, number 6, JUNE 1979.

The Nature of Robots

PART 1: DEFINING BEHAVIOR

A scientific revolution is just around the corner, and anyone with a personal computer can participate in it. The last time this happened, 250 years ago, the equipment was the homebrew telescope and the subject was astronomy. Now, astronomy belongs just as much to amateurs as to professionals. This time the particular subject matter is human nature and in a broader scope, the nature of all living systems. Some ancient and thoroughly accepted principles are going to be overturned, and the whole direction of scientific investigation of life processes will change.

The key concept behind this revolution is control theory. Control theory has been developing for almost 40 years, and has already been proposed (by Norbert Wiener) as a revolutionary concept. It has not been easy, however, to see just how control theory can be made part of existing scientific approaches although many people have tried. Most of these attempts have tried to wedge control theory into existing patterns of thought. To apply any new idea in such a way, while ignoring the new conceptual scheme made possible, is to deny the full potential of the new idea.

Many life scientists who have tried to use control theory have tried to imitate the engineering approach, dealing with human beings as part of a man-machine system instead of complete control systems in their own right. Others have used control theory directly to make models of human and animal behavior, but have concentrated on minor subsystems, failing to see that the organism as a whole can be dealt with in terms of the same principles. The result has often been a strange mixture of concepts—a patchwork instead of a system.

Strangely enough, many engineers who do understand control theory haven't done much better. Here the problem is that these engineers tend to accept the basic concepts developed by biologists and psychologists, and to use control theory to explain cause-effect relationships they are told exist—but which in fact do not exist. We will start this development by looking at something called *behavior*, which biologists and psychologists have assured engineers is very important, thereby leading the engineers astray.

What is all this supposed to mean? A lot is meant, though in different ways. Roboticists, for example, are trying to develop machines which will imitate human organization, and so are the artificial intelligence experimenters. But from whence came the description of the system they are trying to model? Basically, it came from the life sciences. If the life sciences are using the wrong model, it would be essential to know that before much more labor is invested in imitating an imaginary creature.

Perhaps the most general reason control theory is interesting is that it concerns people. There aren't many sciences left in which important discoveries can be made by amateurs working at their own tables. Control theory opens up an entirely new field of experimentation, a kind that has never been done before in psychology or any other life science.

All that is needed by amateurs who want to participate in these developments is a basic grasp of control theory, an understanding of the procedures that go with it, some basic equipment, and curiosity about human nature. I shall now provide the first two items on that list. The rest is up to you.

The Problem With Behavior

The word *behavior* is used frequently—we hear about behavioral science, behavior modification, behavior therapy. For example, *Science News* now has a "Behavior Column"; it was formerly the publication's "Psychology Column". An innocent bystander might conclude that any word this important must have a universally accepted definition, but that is not true. Behavior is a slippery concept.

Here is an example of a person behaving. Chip Chad is seated in front of a teletypewriter pounding keys. What is he doing?

Is he alternately tensing and relaxing muscles in his arms? Yes. Is he moving his fingers up and down? Yes. Is he typing strings of symbols? Yes. Is he adding a return instruction that he forgot at the end of a subroutine? Yes. Is he writing a program for plotting stock market prices? Yes. Is he making a little extra money for a vacation? Yes. Is he justifying his hobby to his family? Yes.

Clearly, each description of what Chip is doing is, in fact, an accurate description of the very same collection of actions. Which one, then, is Chip's behavior? Obviously, they all are expressions of behavior.

Suppose Chip decides that he really doesn't need a subroutine, and substitutes a jump instruction for the return. Now, he is writing the program—obviously the same program—by using a different behavior. Or suppose he buys an input device, and continues working on the subroutine by speaking letters into a microphone. Now he is using different muscles and movements, but he is still doing the same behaviors farther down the list. How could he be doing the same thing by means of doing something different?

Or consider Chip driving a car along a straight road. He is consciously steering. This happens to be a gusty March day, and every five minutes the wind changes speed and direction. Chip is an experienced driver, and continues to steer the car down the road in a straight line. If we look at what his arms are doing, however, we find that they are moving the steering wheel in an apparently random pattern, now centered, now far to the right, now far to the left. Somehow he is managing to produce a constant steering-the-car behavior by means of a behavior that is widely varying. The path of the car doesn't correlate with the position of the steering wheel at all.

Scientists have always thought of behavior as the final product of activity inside the organism. The brain sends commands to the muscles, which create forces, which produce movements, which generate the stable and repeatable patterns we recognize as behavior. There is, in principle, a chain of cause and effect, with the events at the end of the chain being caused by the events at the beginning. Such scientists would say that in the example with Chip at the computer keyboard, we were simply attending to various stages in that chain.

How does that picture fit in with Chip's driving the car in a straight line? The direction in which the car is going is affected by his movements of the steering wheel, and is farther out along the chain of causes and effects. But the wind adds its effects on the direction of the car *after* Chip's effects in the chain. Somehow he is varying his actions so that when their effects are added to the effects of the randomly varied wind, the result is something constant. If we had been thinking of driving the car in a straight line as Chip's behavior, we have to revise that idea: the direction of the car depends just as much on the wind as on Chip.

It may seem that we have simply moved our definition of behavior closer to Chip. But consider *how* he moves the steering wheel. The wheel moves when the forces reflected from the front wheels do not exactly balance the forces created by his muscles. As the car goes along, the roadbed tilts and various bumps and dips cause changes in the reflected forces. The wheel may be turned far to the right, into the crosswind, on the average, but maintaining the wheel in that position requires that his muscles be constantly changing tension, as the reflected steering wheel forces fluctuate. We have the same problem as before: Chip produces a varying output that affects the steering wheel, but the steering wheel is also being affected by forces that are independent of what Chip is doing with his muscles. Yet the *sum* of the muscle forces and those extraneous forces is zero, except when the steering wheel is changing position.

Even if we back up another step and call Chip's muscle tensions his behavior, we have trouble. Muscles are made to contract by signals from the nervous system, but muscles don't respond the same amount to a given signal every time they are used. They fatigue; other muscles interfere with them; joint angles change so that a given muscle tension can produce different amounts and directions of force. The only *behavior* that Chip produces which can be attributed entirely to Chip and not in part to his environment consists of the nerve signals that leave his nervous system and enter his muscles.

If we want to be completely accurate about Chip's behavior, we should consider the output signals from his nervous system, and leave everything else in his environment. That is what we will do, but by doing that we create the biggest problem of all.

A scientist studying a behavior hopes to learn enough about its rules to predict when it will occur. Under the old approach, this means varying factors in the environment and looking for behaviors that correlate with those variations. But if we try to describe behavior in terms of the output signals from the nervous system, all correlations disappear. Oh, maybe we have a knee jerk or a sneeze left over, but we have lost all the regularities that give us some reason to talk about behavior in the first place. We would never guess, from looking at Chip's neural signal outputs, that the result of them would be a straight path of a car that is being forced one way and another by a variable crosswind.

When you pause and reflect upon what has been covered so far, you will realize that we are already deep into control theory, even though we haven't discussed it by name yet. We have dealt with the subject as such because the discussion concerns a fundamental difficulty with the very concept of behavior, especially the concept that behavior is the final product of an organism's inner activities. As we see how this difficulty gets resolved, we will be forced into control theory no matter how we approach the solution. One reason biologists or psychologists have not developed control theory is that they have clung stubbornly to the idea that behavior is part of a causal chain that starts in the nervous system (or in stimuli that cause activity in the nervous system) and propagates outward from there according to physical laws of cause and effect. That is why people design robots in the same way, and why those robots have yet to behave in a way that is convincingly alive. In order to solve this problem instead of just brushing it aside, we have to admit that the causal chain in which people have believed for so long simply does not exist, and never has existed.

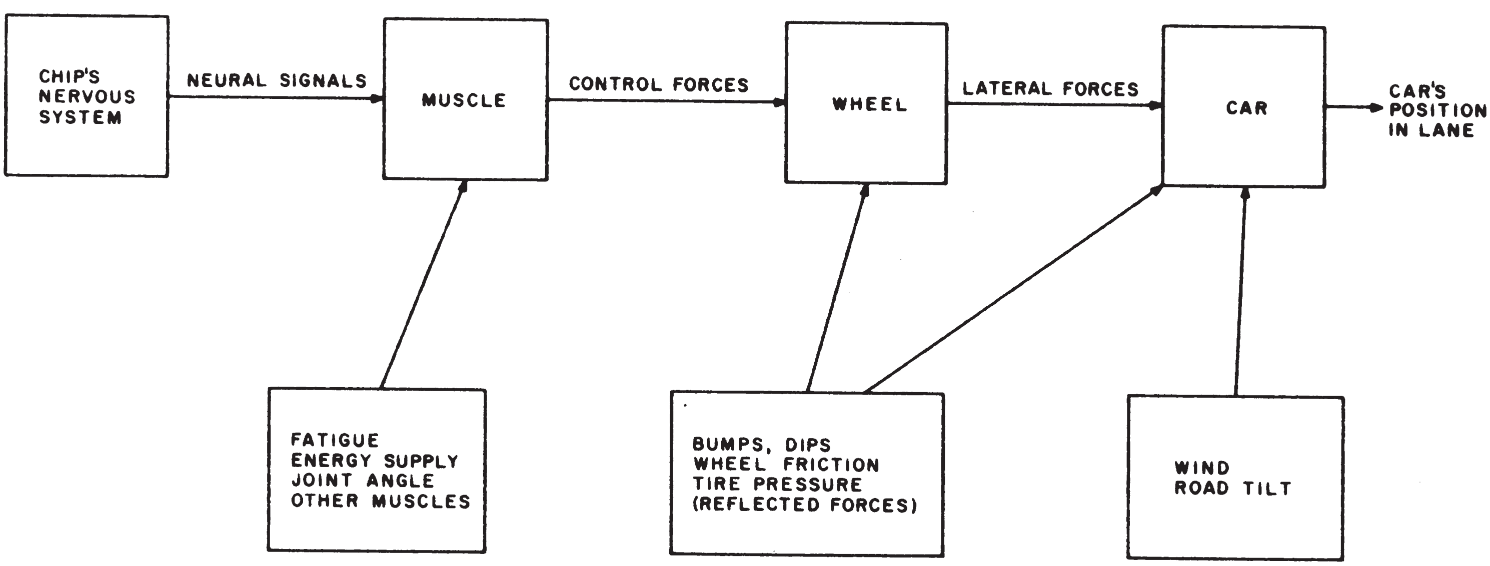

Figure 1: *The cause and effect chain leading to behavior. The behavior called "driving in a straight line" is anything but simple. Some psychologists speak of behavior as simply being emitted by an organism, but this is clearly an inadequate concept. Between the nervous system and the stable pattern it appears to produce, disturbances come into play, having just as much effect on the final outcome as the nervous system has. Nevertheless, the most regularity appears at the end of this chain, and the least at the beginning.*

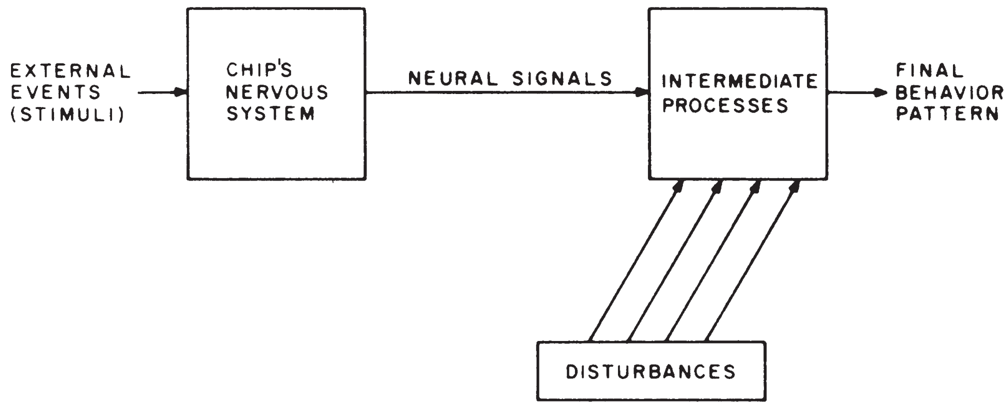

Figure 2: *The old model of behavior. In this old model of behavior, environmental "forces" act on the nervous system to make it produce behavior. The logic of this straight-through, cause and effect chain is spoiled by the presence of disturbances which act after the last physical output of the nervous system (ie: neural signals that activate muscles). This cannot be the correct model for stable behavior.*

Figure 1 sums up the problem we are dealing with. At every stage of events following the outputs from Chip's nervous system *disturbances* come into play, adding to the effects that can be traced to the neural signals. As we go farther to the right of the figure, we might expect that any regularities in Chip's output signals would be lost (ie: that each successive variable would show more and more random variations).

Exactly the opposite is true. The farther to the right we go in figure 1, the less random variation occurs. The variable farthest to the right, the relationship of the car to its lane, can remain constant within a few inches for hour after hour. We find that this is the *most* stable variable in the chain, and that as we go backward up the chain toward Chip's nervous system, the random-looking variations get larger and larger. At the beginning of the chain the variations become totally unpredictable.

Consider figure 2; we added the effects of external events on a nervous system. According to the old picture still fundamental to most life sciences, external events act on the physical structure of the nervous system (along with internal events such as changes in body chemistry), and cause outputs to occur. Those outputs have consequences which show up at the end of the chain as behavioral patterns. To study the organization of behavior, you manipulate the external events, and look for regular behaviors that result (of course, you find them).

But in figure 2 we also see those random disturbances. The only way to get away from them is to make sure that the environment remains absolutely stable (ie: that nothing happens which can interfere with behavior). The standard approach requires eliminating those disturbances, for the simple reason that if they are not eliminated, the experimental results disappear into the background noise. Thus by eliminating disturbances as completely as possible, under the guise of establishing standard (ie: control) experimental conditions, some scientists have swept this basic problem under the rug. They have also done away with the principal tool we have for understanding how these systems really work. If there are no disturbances, then the idea of a cause-effect chain running from external events through the organism to behavior seems to hold up, more or less. As soon as natural disturbances are allowed to occur, we find that the *overall* connection from external event to final behavior remains as clear as ever; but, the model of what happens in between falls to pieces with a loud crash.

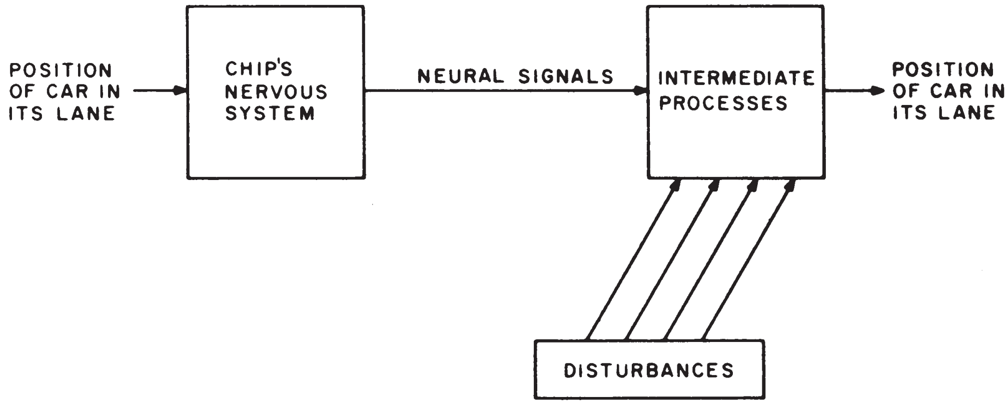

Figure 3: *A slightly different view of the old model of behavior. The principle stimulus involved in driving a car in a straight line is the position of the car in its lane. This is the same variable that is the measure of behavior. The variable that is the final outcome of Chip's actions is the same variable that provides inputs to the nervous system that is acting. The variable at the causal end of the chain is the same variable found at the effective end of the chain.*

Closing the Loop

There seems to be nothing wrong with figure 2; nothing, that is, except that it cannot account for the regularities of behavior. There is something wrong; something has been left out. Let's focus on the final variable in the chain, the position of the car relative to the lane. What variable that could affect Chip's senses, do you suppose, would have the most to do with his manipulations of the steering wheel? The position of the car relative to the lane. This variable is both the consequence of Chip's actions, and the main source of sensory information that could cause him to act (see figure 3).

Psychologists have gone this way before. They have tried to make sense of this situation by supposing that the behavioral variable is somehow different from the stimulus variable. If the position of the car relative to its lane is the behavioral variable, then perhaps the onset of a change in the visual image of the road is the stimulus variable. That leads to the idea of a *chain* of stimuli and responses. The car drifts in its lane; that stimulates Chip's nervous system to make a response, which affects the physical position of the car in its lane, which causes a new change in the stimulus, and so on around and around.

There are several severe difficulties with this explanation. In the first place, there is no way to separate the visual image from the position of the car; these are just two ways of talking about one whole physical situation in which a certain collection of interdependent variables changes simultaneously. The alternation between stimulus and response is completely imaginary, as anyone who drives knows. If causes and effects really were sequential, and chased themselves around and around the loop, it is unlikely that Chip would keep the car on the road for more than ten seconds. I n part 2 we'll do a proper simulation in BASIC, and you will see that when the system is designed to behave sequentially, the result is most likely to be violent oscillations.

There is no reason at all to make an artificial distinction between the position of the car on the road as a behavioral response and as the stimulus which causes the response. Only one physical situation exists, and there is no need to present it in two disguises. The position of the car on the road is both an effect of Chip's actions and the sensory situation which leads (with a little help from Chip) to those actions. There is a closed loop of cause and effect, and the position of the car is just one part of that loop.

Now we begin to draw a diagram of a proper control system. In figure 4, three physical quantities are shown, an *output* quantity, an *input* quantity, and a *disturbing* quantity.

The output quantity corresponds to an output of Chip's that is entirely due to himself (ie: perhaps due to the neural signals reaching his muscles or to some variable farther down the chain of figure 2, revealed when disturbances are known or can be legitimately eliminated).

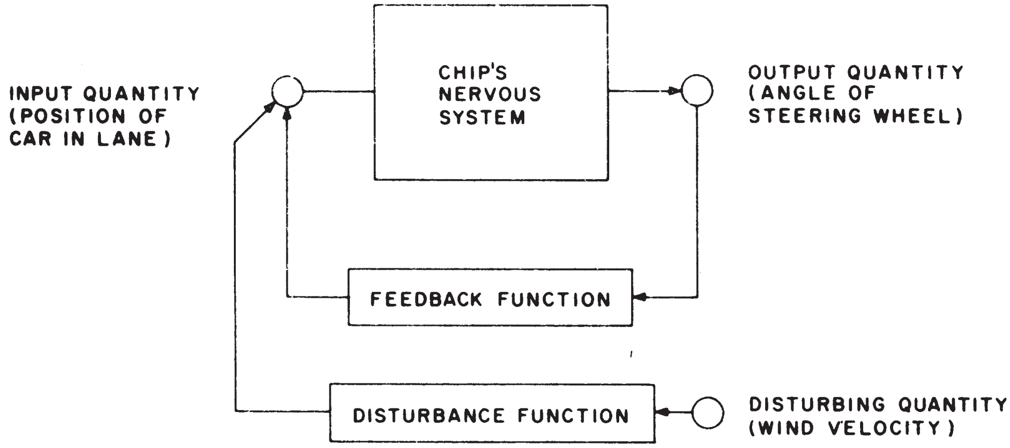

Figure 4: *Closing the loop. By rearranging the relationships shown in figure 3 and eliminating the redundant appearance of the car position, we create a closed loop diagram. This is the general form of a control system diagram that will be used in this series from now on. The controlled variable is always the input quantity; the output quantity is the means of control. The single disturbance shown represents the net effective disturbance if more than one is acting at the same time. The disturbing function is chosen to provide the proper net contribution to the input quantity. The feedback function represents links external to the behaving nervous system through which outputs are transformed into contributions to the state of the input quantity.*

The input quantity is the variable that is stabilized by the variations in Chip's output. Thus we call the input quantity, here, the position of the car relative to its lane. Of course, by that we mean whatever it is about that position that can be a sensory input to Chip (ie: probably a visual image of the hood of the car and the road beyond, framed in the windshield).

Between the output quantity and the input quantity is placed a *feedback function*. This function expresses the physical links that exist between Chip's output quantity and the input quantity. In the case of a moving car, if the output quantity were the angle of the steering wheel, which it might be if the angle is also a controlled quantity, then the effect of the wheel angle would be a continual change of car position, and the feedback function would have to include at least one time integration. The feedback function is simply a description of the physical processes which give each magnitude and direction of the output quantity a contribution to the state of the input quantity.

In figure 4 we also include disturbances as an integral part of the diagram of the system. The disturbing quantity in this case would be wind velocity and direction, and the *disturbance function* connecting it to the input quantity would express the way in which aerodynamic laws convert wind velocity into effects on the car's position in its lane.

The state of the input quantity, therefore, can be expressed in terms of all effects which contribute to it. We have shown only the output quantity and the disturbance due to wind. Many other disturbances—low tires, or tight wheel bearings, or gradation in the road—could also contribute to the state of the input quantity at the same time. All disturbances, however, can be reduced to a single one, since no matter what the cause of the disturbance, the only effect that matters is the effect on lateral position of the car.

Chip himself can be represented by a function, a function that converts the sensed position of the car into a steering wheel angle. This *system function* (system, being short for behaving system) will surely contain delays, nonlinearities, and even variations of its parameters. At first glance it may seem a terrible oversimplification to reduce a whole human being to a simple input/output box, but the situation isn't that bad. We are centering this diagram around the input quantity, not around Chip as a whole; therefore the "Chip box" does not wholly represent him, but only that part which reacts to changes in the input quantity by altering the output quantity. Furthermore, the Chip box (ie: the system function) is not quite as simple as it seems even after being simplified a great deal.

The functions connecting the variables in this closed loop can be extremely complex, and even to approach this system analytically will obviously require some approximations. This is not the place to justify every simplification; sometimes complex mathematics are required to reach a simple conclusion. I'll drop some hints along the way about how the simplified model is generated and why it works, but if you really want to get into this, study a text on servomechanism design.

Simulating Chip

Let us conclude by building a working simulator of Chip driving the car. This is just a hint of what this 4 part series of articles will develop. Building the simulator requires building some special numbers into the program without any explanation at present. The point is to enjoy the simulation, and get used to the idea that everything in a control loop happens at the same time.

We will assume that the steering wheel angle to left or right of center is Chip's output quantity, and that there are no disturbances that can interfere at this point. This output quantity will be called A.

Under the influence of A alone, the car would drift sideways at a rate proportional to A, for small deviations from the center of the lane. Designating the crosswind velocity as W, if W were the only influence acting, the car would drift sideways at a rate proportional to W (in this somewhat oversimplified universe). In the BASIC program we will assume that each iteration corresponds to a fixed amount of elapsed time, so the distance D that the car will drift during any one iteration is simply the sum of the two influences acting on it (line numbers correlate with listing 1):

```
7   D = K1 * W + K2 * A
```

The position, I, of the car relative to its lane will change by an amount D on each iteration:

```
8   I = I + D
```

Now I must introduce a detail: if we just had Chip respond proportionally to the deviation of car position, we would have to make his muscles so flabby that hardly any response would occur, unless we wanted to demonstrate selfimmolating oscillations. We have to take care of two destabilizing factors. First, the feedback function is essentially an integrator, and so puts a lag into the control process. This alone would not cause a problem, but Chip also contains a *transport lag*; he cannot actually produce an output at' the same instant that the input occurs, nor can our program since it is evaluating equations one at a time. The integration lag we take care of by adding to the position I (which Chip senses) the variable D, which is approximately the first derivative of the input quantity. He senses the input quantity with some emphasis on its rate of change, which is actually a realistic model of human perception. This part of the stabilizing of the control action is done in step 9:

```
9   A1 = K3 * (I + 0.8 * D)
```

We have computed a variable Al, the angle which the wheel would assume if Chip reacted instantly. But to handle the transport lag, we must slow his response, letting only a fraction K5 (between 0 and 1) of it occur during any one iteration. That is what step 10 does:

```
10   A = A + K5 * (A1 – A)
```

This slowing technique will be used in the larger simulator next time. To see how it works, set Al to 10.00, K5 to 0.25, and A to 0, and then simply keep doing step 10 with pencil and paper. A will gradually approach the value of Al from any starting point.

The program in listing 1 asks for a wind velocity, and then proceeds to do ten iterations of the control loop, printing wheel angle A and car position deviation I each time. A positive number means the wind is blowing, the wheel is cocked, or the car has moved to the *right*. If you want to follow the program for more than ten iterations, give it the same wind again. It always starts where it left off.

In part 2, we will begin exploring a model of the kind described in figure 4 and start the somewhat mind boggling task of retraining the intuition to think in closed loop terms instead of straight through cause and effect. There is a *big* difference. We'll see that, in general, control systems control what they sense, not what they do. We'll discover something called a *reference signal*, which functions in a control system exactly the way an inner purpose has always been supposed to function. In part 2, we'll see how perception figures into control. And we'll start working with a more extended BASIC simulator than the tiny one in listing 1. Parts of this simulator will be suitable for building into the computer part of a robot, should anyone want to carry matters that far. ■

```
1    INPUT "WIND, MPH: ",W
2    PRINT "WHEEL ANGLE, DEGREES",
3    PRINT TAB(25),"CAR DEVIATION, FEET"
4    FOR J=1 TO 10
5    PRINT %7F1,"  ",A*10,
6    PRINT %7F1,TAB(25),I
7    D=.05*W+A
8    I=I+D
9    A1= –2*(I+.8*D)
10   A=A+.200*(A1–A)
11   NEXT
12   GOTO 1
13   END

RUN

            WIND, MPH: 20
      WHEEL ANGLE, DEGREES        CAR DEVIATION, FEET
               .0                         .0
             –7.2                        1.0
            –11.8                        1.3
            –13.3                        1.1
            –12.7                         .8
            –11.3                         .5
            –10.1                         .4
             –9.5                         .4
             –9.4                         .4
             –9.6                         .5
            WIND, MPH: –30
      WHEEL ANGLE, DEGREES        CAR DEVIATION, FEET
             –9.8                         .5
              8.0                       –2.0
             19.3                       –2.7
             23.1                       –2.2
             21.6                       –1.4
             18.3                        –.8
             15.4                        –.4
             13.8                        –.4
             13.6                        –.5
             14.0                        –.7
            WIND, MPH: 40
      WHEEL ANGLE, DEGREES        CAR DEVIATION, FEET
             14.6                        –.8
            –10.1                        2.7
            –26.0                        3.7
            –31.2                        3.1
            –29.2                        2.0
            –24.6                        1.0
            –20.5                         .6
            –18.4                         .5
            –18.0                         .7
            –18.6                         .9
            WIND, MPH: –50
      WHEEL ANGLE, DEGREES        CAR DEVIATION, FEET
            –19.5                        1.0
             12.3                       –3.4
             32.6                       –4.7
             39.3                       –3.9
             36.8                       –2.5
             30.9                       –1.3
             25.7                        –.7
             22.9                        –.6
             22.4                        –.9
             23.2                       –1.1
```

Listing 1: *A rough simulation of Chip driving the car in a straight line. Each iteration is assumed to correspond to a fixed time interval. Therefore, the distance the car drifts away from straight line travel is the sum of the wind and steering wheel angle. The simulation shows Chip trying to arrive at the wheel angle which will counteract the force of the blowing wind. If you repetitively use the same wind value, you will see that a steady wheel angle is arrived at. [I found it interesting that this simulation seems to settle down within 60 time units to a consistent value. Even changing wind values from +1000 to –1000 units was compensated for within 60 time units....* **RGAC***]*