Journal Title: Behavior research methods, instruments & computers : a journal of the Psychonomic Society, Inc.

Article Author: Psychonomic Society. Marken, Richard S.,

Article Title: Marken, Richard S.,: Spreadsheet analysis of a hierarchical control system model of behavior.

Volume: 22 **Issue:** 4

Month/Year: 1990 **Pages:** 349-359

IUA TN: 688394
ILL Number: 167160276

EMAIL: WEDELIVER@ANTIOCH.EDU

Behavior Research Methods, Instruments, & Computers
1990, 22 (4), 349-359

— METHODS & DESIGNS —

Spreadsheet analysis of a hierarchical control system model of behavior

RICHARD S. MARKEN
Aerospace Corporation, Los Angeles, California

The behavior of a hierarchy of control systems can be simulated with an electronic spreadsheet. Each control system is a column of three cells representing the reference signal, perceptual signal, and output variable of the system. All of the control systems are closed-loop, the input to each system being a function of its output. The circular references in the spreadsheet are resolved through iterative recalculation. When the parameters of each control system (amplification and slowing factors) are set to appropriate values, all control systems in the hierarchy continue to match perceptual signals with reference signals. A three-level hierarchy with four systems at each level is simulated in the spreadsheet. The spreadsheet model makes it possible to observe the dynamic behavior of the control systems as they correct for the effects of environmental disturbance and changes in higher order reference signals. It is possible to "reorganize" the system by changing the perceptions controlled by systems at different levels of the hierarchy. The user can also test to determine the variables being controlled by the system.

Powers (1973, 1989) has proposed a hierarchical control system model of the purposive behavior of organisms. In this paper, a method of implementing the model on an electronic spreadsheet will be described. The spreadsheet system makes it relatively easy to explore the model's behavior with a personal computer. A spreadsheet implementation was selected because many (perhaps most) personal computer users have access to spreadsheet software and are familiar with its use. Moreover, the matrix design of the spreadsheet provides an excellent format for representing a control system hierarchy. The spreadsheet model is designed to be a self-instructional system for those who want to learn how control systems work, but it could also be used as a research tool. The behavior of an appropriately designed version of the model could be compared to that of human or animal subjects in experiments like those described by Marken (1986, 1989) and Bourbon (1989).

HIERARCHICAL CONTROL SYSTEM MODEL

Basic Control System

The components of a basic control system are shown in Figure 1. The *sensor* converts an input variable, *i*, into a perceptual signal, *p*. The *comparator* subtracts the perceptual signal from a reference signal, *r*, to produce an error signal, *e*. The *amplifier* converts the error signal into an output variable, *o*. (Note: *Signals* are quantities that vary inside the control system; *variables* are quantities that vary outside the control system.) What constitutes an input and an output variable depends on the location of this basic system in a control hierarchy. If the system is at the lowest level of the hierarchy (as is shown in Figure 1), then input and output are physical variables in the environment. If the control system is higher in the hierarchy, then input and output are signals coming from and going to lower level control systems; the lower level systems are the "environment" of the higher level systems.

Regardless of their position in the hierarchy, all control systems are designed to do the same thing—keep the input variable, *i*, in a predetermined state specified by the reference signal, *r*. The problem of control arises because the value of the input variable is affected by system outputs as well as disturbances, *d*. A disturbance is any external influence on the input variable that is not caused by the system itself. When set up properly, a control system produces outputs that counteract the effects of disturbances on the input. The input variable, which is maintained at a value that corresponds to that specified by the reference signal, is called the *controlled variable*. The value of the input that corresponds to the setting of the reference signal is the *reference state* of the controlled variable. The reference state is constant if the reference signal is constant, and it varies if the reference signal varies. However, constant or varying, the controlled variable is kept in the reference state, continuously protected from the effects of disturbance by the output of the control system.

Correspondence should be addressed to Richard S. Marken, Aerospace Corporation, Mail Station M1/076, P.O. Box 92957, Los Angeles, CA 90009.

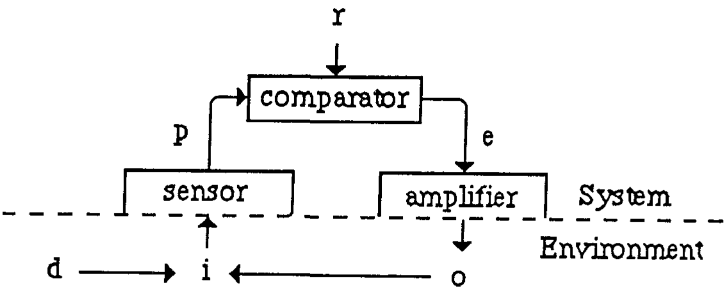

Figure 1. A basic control system.

A concrete example of a basic control system is the light intensity control system of the eye (sometimes called the "pupillary reflex"). The input variable, *i*, is the light intensity and the output variable, *o*, is the size of the pupil. The sensor (retina) converts light intensity into a perceptual signal. The comparator converts the difference between the perceptual signal and a reference signal (from higher levels of the brain) into an error signal. The amplifier converts the error signal into muscle tensions, which increase or decrease the size of the pupil (the output variable). The size of the pupil affects the intensity of the light incident on the retina, so that the output of the system affects the input. Since the input also affects the output, via the control system, there is a closed loop of cause and effect. The intensity of light at the retina is the controlled variable. Light intensity will be kept in some reference state, despite disturbances such as variations in the intensity of the light source. Disturbances are precisely countered by variations in pupil size (output), which can be seen if you vary the intensity of a light while looking in the mirror.

Higher Level Systems

In a hierarchy of control systems, the lowest level systems are connected to the external environment; input and output are physical variables. Higher level control systems receive inputs from and send outputs to these lower level systems. Inputs to higher level systems are perceptual signals from the lower level systems, and outputs from higher level systems become the reference signals of the lower level systems. The relationship between three levels of control in a control hierarchy is shown in Figure 2. Note that systems at all levels are closed-loop; there is a connection, through the environment, from output to input and, through the control system, from input to output. The outputs of higher level systems pass through one or more layers of lower order systems before entering the environment. The inputs to higher level systems pass from the environment through one or more layers of perceptual processing before becoming higher order perceptual signals.

The sensor component of a higher level control system transforms one or more perceptual inputs from lower level systems into a single perceptual signal. The nature of the transformation can, in principle, be quite complex, making the perceptual signal a measure of variation in some abstract aspect of the environment. For example, the sensor might compute a perceptual signal that is a weighted sum of several lower level perceptual signals. This weighted sum is the variable that is controlled by the system. The reference signal specifies the reference state for this weighted sum (actually, for the perceptual signal that represents the weighted sum). The system controls this perceptual signal by varying its outputs, which become the reference signals of lower level systems. These reference signals tell the lower order control systems what to perceive, not what to do; reference signals are specifications for input, not output. The higher level systems control their inputs by specifying the level of input to be perceived by lower order systems.

A Working Model

The hierarchical control model is designed to produce purposive behavior like that seen in living organisms. This behavior is difficult to visualize in a static representation of the model, like that in Figure 2. It is not obvious that control systems at all levels of the hierarchy can achieve their goals simultaneously, even in the presence of randomly varying disturbances. The static representation can even lead to misconceptions about the capabilities of a control hierarchy. Fowler and Turvey (1978), for example, claimed that a two-level control hierarchy, similar to the one in Figure 2, could not achieve two different goals simultaneously. Unfortunately, mistakes like this, made by authoritative authors, have led other researchers to reject control system models before their capabilities have been explored. The spreadsheet model described in this paper was developed, in part, to remedy this problem.

There is a large literature on the theory of control systems, but it can be difficult to understand the behavior of these systems by looking at the equations that describe them. Those who want to learn the capabilities of a control system model of behavior, especially those who are not mathematically sophisticated, would benefit from a working, dynamic simulation of control system behavior. A spreadsheet implementation of this simulation has several advantages over conventional languages. First, the matrix layout of the spreadsheet is well suited for the display of a hierarchical control model. Second, the modular design of the spreadsheet makes it relatively easy to change the model (adding levels or adding systems to existing levels).

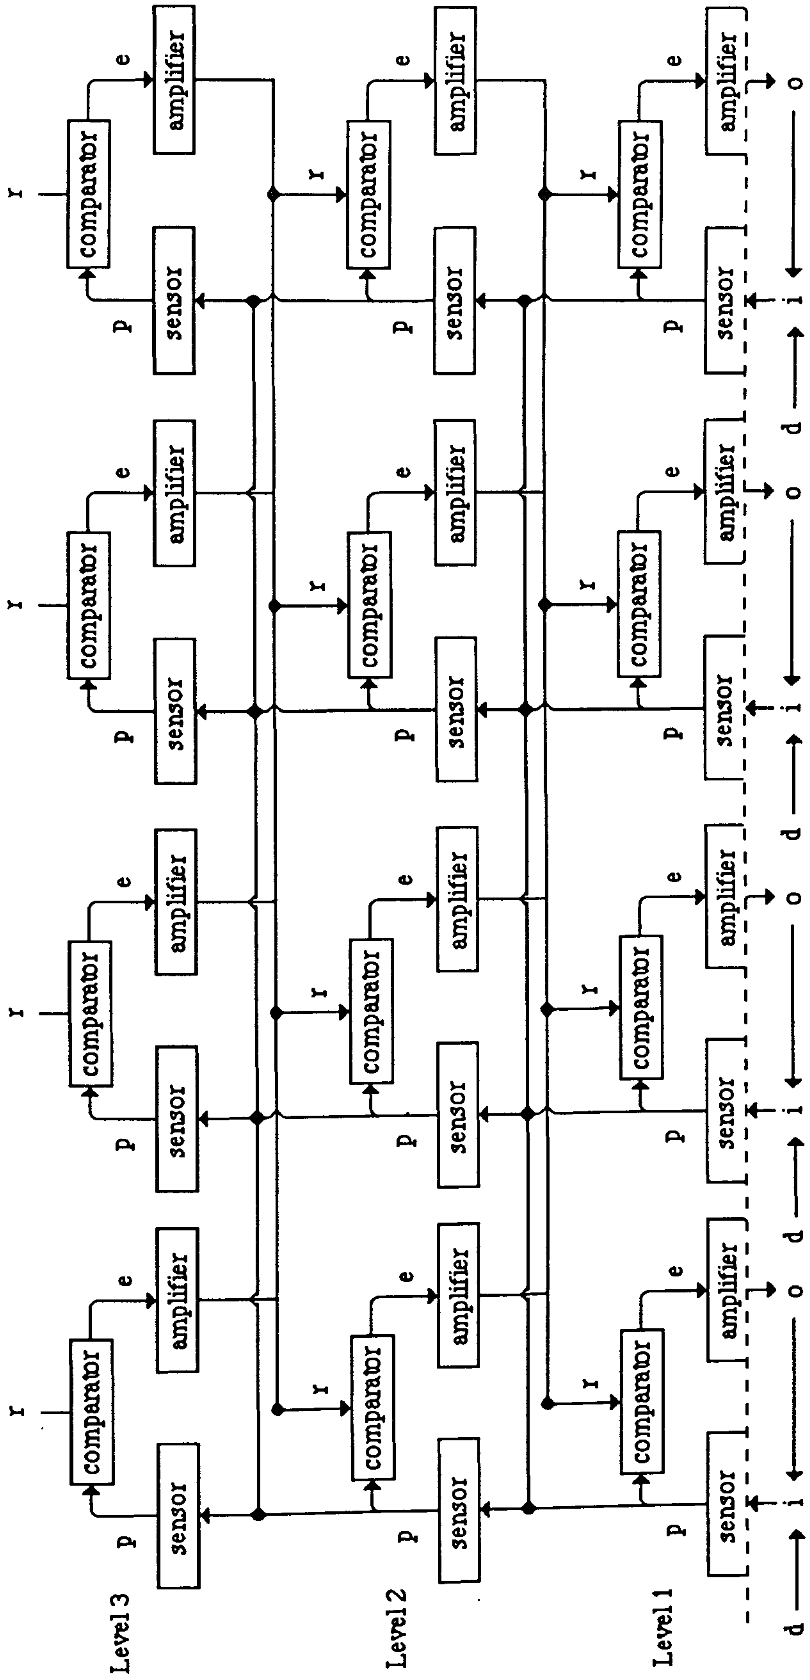

Figure 2. A three-level hierarchy of control systems.

Third, it is easy to change model parameters by typing a new value into an appropriate cell of the spreadsheet. Finally, the spreadsheet is a "programming environment" that is familiar to most users of personal computers.

THE SPREADSHEET CONTROL HIERARCHY

A control system can be represented in an electronic spreadsheet as a column of three cells, as is shown in Figure 3. Each cell contains a formula that computes the value of a signal or variable in the control system. The top, middle, and bottom cells contain formulas for the reference signal, perceptual signal, and output variable, respectively. The formulas for the reference and perceptual signals differ, depending on the level of the control system in the control hierarchy. However, the computation of the output variable is the same for all control systems.

Output Signal Calculation

The output signal is proportional to the difference between the reference and perceptual signal. Algebraically, the formula for the output signal can be written as a difference equation as follows:

$$o(t+1) = o(t)+s[g(r-p)-o(t)],$$

where o is the output signal at time t; s and g are constants representing slowing and gain factors, respectively; r is the reference signal; and p is the perceptual signal. The equivalent spreadsheet version of the formula, written in Lotus 1·2·3 (LeBlond & Cobb, 1985) is

+O+SLOW*(GAIN * (R−P)−O).

The spreadsheet formula is contained in Cell O, making it self-referential.[1] The amplified error signal is accumulated into Cell O (output). The error signal is the difference between reference, R, and perceptual, P, signals (R−P). The amplification factor is in the cell called GAIN. The rate at which the amplified error is accumulated into O is determined by the slowing factor in the cell named SLOW. The accumulated output in Cell O is equivalent to the time integral of the error signal. The integration performed in real nervous systems is not perfect; there is some loss or "leakage" from the integrator over time. This leakage is captured by the loss due to slowing in the equation for the output function.

The appropriate values of GAIN and SLOW depend on the level of the basic control system in the hierarchy. SLOW determines the "speed of response" of the control system. In a hierarchy of control systems, higher level systems cannot respond more rapidly than lower level systems. This means that the value of SLOW for higher order systems must be the same as or smaller than that for lower level systems; the higher level systems should accumulate output more slowly than lower level systems do.

The GAIN of the system determines its sensitivity to error; the higher the value in GAIN, the greater the output per unit error. High GAIN means precise control, but systems with high GAIN must respond more slowly than systems with low GAIN. The values of GAIN and SLOW must be inversely related.

Perceptual Signal Calculation

Powers (1973) has suggested that the perceptual signals at different levels of a control hierarchy represent different classes of perceptual variables. The lowest level perceptual signals represent only the *intensity* of the input. This level of the hierarchy deals with the outside world only in terms of magnitude of the input incident on the sensors. The next level of perception is *sensation*—a function of several intensities. Going up from sensations, there are *configurations* (combinations of sensations), *transitions* (temporal changes in configurations), *events* (a sequence of changing configurations), *relationships* (perceived relationship—logical, statistical, causal, etc.—between separate events), *categories*, *programs*, *principles*, and *system concepts* (see Powers, 1989, pp. 190–208). In the spreadsheet model, there are only three levels of perception—intensity (Level 1), sensation (Level 2), and relationship (Level 3).

Level 1. The formula for the Level 1 (intensity) perceptual signal is $p = i$, where p is the perceptual signal and i is an environmental input variable. The spreadsheet version of this formula is just +I in the perceptual signal cell of the Level 1 control system. The formula represents the sensor as a perfect linear transducer that transforms an input variable, located in Cell I, into a perceptual signal.

Setting the Level 1 perceptual signal to the value in Cell I implies that the perceptual signal is an exact measure of the physical variable. A more realistic model of the sensor might represent the perceptual signal as a logarithmic or exponential function of the input variable. This can be done by using a spreadsheet function, such as @LOG(I), to compute the Level 1 perceptual signal. These equations must take account of the fact that I can

	Cell Values	Cell Formulas
R(j,i)	7	+E4-D4
P(j,i)	6.004	@INDEX(PW,1,PW21)*P11+@INDEX(PW,2,PW21)*P12...
O(j,i)	41.63	+O21+SLOW*(GAIN*(R21-P21)-O21)

Figure 3. The basic control system as implemented in the spreadsheet.

PW Matrix

Row Labels				
1	1	1	1	1
2	1	1	1	-1
3	1	1	-1	1
4	1	1	-1	-1
5	1	-1	1	1
6	1	-1	1	-1
7	1	-1	-1	1
8	1	-1	-1	-1
9	-1	1	1	1
10	-1	1	1	-1
11	-1	1	-1	1
12	-1	1	-1	-1
13	-1	-1	1	1
14	-1	-1	1	-1
15	-1	-1	-1	1
16	-1	-1	-1	-1

Figure 4. The perceptual weighting matrix, PW.

be negative. The model has been successfully run using a logarithmic transform:

$$p = \log(i) \quad i > 0,$$

where p is the perceptual signal and i is the input variable. The corresponding spreadsheet formula is:

@IF(I>0,@LOG(I)*10,−10).

The log of I is multiplied by 10 to increase the dynamic range of the perceptual signal. If the physical signal is less than or equal to 0, the perceptual signal becomes an arbitrary (and fairly large) negative number, −10.

Level 2. The Level 2 perceptual signal is a weighted sum of four Level 1 perceptions. The weights are selected from a set of 16 in a matrix called PW (an abbreviation for *perceptual weights*). The PW matrix is shown in Figure 4. The weights assigned to a Level 1 perception are either 1.0 or −1.0, representing an excitatory or inhibitory connection of the Level 1 perceptual signal to the Level 2 sensor. Each Level 2 sensor uses weights from a different row of the PW matrix. The numbers identifying the rows of PW that are used as weights by Level 2, Systems 1 through 4, are stored in pointer cells named PW21 through PW24, respectively. PW21 contains a pointer (a number between 1 and 16) to the row of the PW matrix that contains the weights used by Level 2, System 1; PW22 contains a pointer to the row of the PW matrix that contains the weights used by Level 2, System 2, and so forth.

The following formula was used to compute the perceptual signal for Level 2, System 1:

$$p = \sum_i w_{1i} p_i,$$

where the sum is over the $i = 1$ to 4 Level 1 systems, w_{1i} is the weight assigned to Level 1, Perceptual Signal i by Level 2, System 1, and p_i is Level 1, Perceptual Signal i. The corresponding spreadsheet formula for this perceptual signal is

```
@INDEX(PW,1,PW21)*P11 +
@INDEX(PW,2,PW21)*P12 +
@INDEX(PW,3,PW21)*P13 +
@INDEX(PW,4,PW21)*P14.
```

The formula uses the @INDEX function to access the appropriate row of weights in the PW matrix. Cell P11 is the perceptual signal from Level 1, System 1; Cell P12 is the perceptual signal from Level 1, System 2, and so on. Cell PW21 contains a number corresponding to the row in PW to be used in computing the perceptual signal for Level 2, System 1. The statement @INDEX(PW,1,PW21)*P11 multiplies Level 1, Perceptual Signal 1 by the perceptual weight in row PW21, column 1 of Matrix PW. If the value in PW21 is 3, then the weight comes from column 1 of the third row of Matrix PW. By changing the values in Cell PW21, it is possible to change the perceptual weights (and, thus, the nature of the perception) computed by this Level 2 system.

Level 3. The Level 3 perceptual signals are logical functions of one or more Level 2 perceptions. The logical expression for the perceptual signal computed by Level 3, System 1 is

$$\text{if } p_{21} > p_{22} \text{ then } p_{31} = 1 \text{ else } p_{31} = -1,$$

where p_{21} and p_{22} are the perceptual signals of Level 2, Systems 1 and 2, respectively, and p_{31} is the perceptual signal of Level 3, System 1. The spreadsheet representation of this formula is

@IF (P21>P22,1,−1),

and it is located in Cell P31. P21 and P22 are perceptual signals computed by Level 2, Systems 1 and 2, respectively. The value of the function is 1 if P21 is greater than P22; otherwise, it is −1. Other systems at Level 3 also use the @IF function to determine whether or not other relationships (>, <, =, < >) hold between various Level 2 perceptions.

Reference Signal Calculation

Outputs from several higher level systems are combined to form the reference signal for a single lower level system. The contribution of a particular higher level system to the reference signal of a lower level system depends on how that lower level system's perceptual signal contributes to the perceptual signal of the higher level system (Powers, 1979a). The basic rule is that the sign of the higher level system's contribution to the lower level system's reference signal must be the same as the sign of the lower level system's contribution to the higher level system's perceptual signal.

Level 1. The reference signal for each Level 1 system is a weighted combination of the outputs from the Level 2 systems. The @INDEX function is again used to determine the weights assigned to the higher level system's contribution to each Level 1 reference signal. The following formula was used to compute the reference signal for Level 1, System 1:

$$p = \sum_i w_{i1}\, o_i,$$

where the sum is over the i = 1 to 4 Level 2 systems, w_{i1} is the weight assigned to Level 1, Perceptual Signal 1 by Level 2, System i, and o_i is the Level 2 output variable from System i. The corresponding spreadsheet formula for this reference signal is

@INDEX(PW,1,PW21)*O21 +
@INDEX(PW,1,PW22)*O22
@INDEX(PW,1,PW23)*O23 +
@INDEX(PW,1,PW24)*O24.

The cells O21 through O24 are the outputs of Level 2, Systems 1 through 4, respectively. Cells PW21 through PW24 contain the PW row indexes of the weights given to the Level 1, System 1 perceptual signal by Level 2, Systems 1 through 4, respectively. The reference signals for other Level 1 systems use the same formula; only the second argument of the @INDEX function (1 in this case) is changed—to 2 for System 2, 3 for System 3, and 4 for System 4.

Level 2. The reference signals for the Level 2 systems depend on how the Level 2 perceptual signal was used in the computation of the Level 3 perceptual signal. Level 3 perceptual signals represent relationships between Level 2 perceptions. In order to preserve negative feedback, the following rule was used to determine how the Level 3 output variables contribute to the Level 3 reference signals: If an increase in the Level 2 perception produces an increase in the Level 3 perception, then the sign of the Level 3 output connection to the Level 2 reference signal is positive; if an increase in the Level 2 perception produces a decrease in the Level 3 perception, then the sign of the Level 3 output connection to the Level 2 reference signal is negative. For example, if the Level 3 perception is @IF(P21 > P22,1,−1), an increase in perception P21 (Level 2, System 1) will tend to make the Level 3 perception increase (toward 1). Thus, the Level 3 output contribution to the Level 2, System 1 reference signal is positive.

Level 3. The Level 3 reference signal values are entered as constants by the user. In a real system, these signals would be genetically determined (if the hierarchy had only three levels) or provided by even higher level systems that perceive variables that are functions of the Level 3 perceptual signals. Any value can be entered for the Level 3 reference signals, but, because of the way in which the Level 3 perceptual signals are computed, it only makes sense to enter values of 1 or −1.

The Complete Hierarchy

The complete spreadsheet model of a control system hierarchy is shown in Figure 5. There are three levels of control systems with four systems at each level. The first two columns of the display contain the values of GAIN and SLOW for systems at each level. Note that the higher level systems (toward the top of the spreadsheet) have smaller values for SLOW than the lower level systems, making the higher level systems respond more slowly than lower level systems. Because the higher level systems respond more slowly, their GAIN can be larger.

The third column of the spreadsheet contains the labels of the signals and variables in control systems at each level of the hierarchy. R, P, and O identify the reference signal, perceptual signal, and output variable, respectively. Numbers in parentheses identify the level and system number of each signal and variable. Thus, R(2,i) identifies the reference signal for Level 2, System i (the system identification numbers, i, which go from 1 to 4, are found in the top row of the spreadsheet). The numbers in the four cells to the right of each identifying label are the momentary values of the signals and variables in the four control systems at each of the three levels of the hierarchy. The column labeled "average error" gives the average value of the error signal at each level of the hierarchy. Each average is taken over the four error signals at one level of the hierarchy.

A single horizontal line separates the control system hierarchy from the physical environment. The first row of four numbers below this line (labeled "input variable") represent stimulation at the sensors of the Level 1 control systems. This stimulation is caused by physical variables that are outside the control hierarchy—the row labeled "disturbances"—as well as by actions of the control hierarchy itself (the Level 1 outputs). The formula for the input variable for Level 1, System i is

$$i_{1i} = o_{1i} + d_i,$$

where i_{1i} is the input to Level 1, System i, o_{1i} is the out-

		System (i)	1	2	3	4		
Delay	Gain	R(3,i)	-1	1	1	1	Average	
Level 3		P(3,i)	-1	1	1	1	Error	0.00
1E-05	500	O(3,i)	-11.6	38.58	58.58	74.21		
		R(2,i)	-11.6	50.19	20	15.63	Average	
Level 2		P(2,i)	-11.8	49.97	20.17	15.63	Error	0.14
1E-04	150	O(2,i)	21.86	17.98	0.53	21.25		
		R(1,i)	3.88	19.12	17.9	-16.8	Average	
Level 1		P(1,i)	3.87	19.1	17.9	-16.8	Error	0.01
0.001	70	O(1,i)	53.87	-69.1	67.89	-33.2		
Input Variable		I	3.87	19.11	17.89	-16.8		
Disturbance		D	-50	50	-50	50		
Test Variable	12	Weights		-1	1	-1	-1	
Stability Factor		327.2				Behavior		19.49

Figure 5. Display of a three-level hierarchy of control systems implemented in the spreadsheet.

put of that system, and d_i is the disturbance to Input i. The spreadsheet version of this formula, contained in the appropriate input variable cell, is +O+D. The input variable is a result of the combined effects of a disturbance, D, and a Level 1 system output, O. For example, the input variable might be the intensity of sound at the sensory surface. The value of this variable depends on the intensity of sound sources in the environment (D) as well as on the actions of the behaving systems (movements relative to the sound sources). A proximal input can be influenced by more than one disturbance and by more than one system output. The spreadsheet model will work in these cases as long as at least one of the outputs affecting the input variable comes from the system controlling that input.

The part of the spreadsheet below the row of disturbance values is used for testing the model. The use of the "test variable," "weights," "stability factor," and "behavior" cells is explained below.

It should be noted that this model is not meant to represent behavior in a particular situation. The goal is to show that a working hierarchical model can be constructed. Once the basic principles of the hierarchical model are grasped, it should be possible to develop a hierarchical model of some specific behavior, such as "pointing at a target" or "lifting a glass."

RUNNING THE MODEL

The following description of the behavior of the model is best understood if the reader follows along with the spreadsheet model up and running.[2] The completed hierarchical control model consists of the equations defining the reference signal, perceptual signal, and output variable for each of the four control systems at each level of the hierarchy. The numerical values for GAIN and SLOW, representing gain and slowing factors for each level of the model, must also be set. Some values that produce stable behavior are shown in Figure 5.

Changing Goals

The values of the highest level reference signals (Level 3) are the ultimate "goals" that the hierarchy of control systems is working toward; all the actions of the hierarchy are aimed at producing the perceptions specified by these Level 3 references. Figure 5 shows these reference values set to −1, 1, 1, and 1. Because of the way in which the relationship perceptions are defined, Level 3, System 1 (with a reference signal value of −1) will try to keep p_{21} less than or equal to p_{22}; Level 3, System 2 will try to keep p_{22} greater than p_{23}; Level 3, System 3 will try to keep p_{23} greater than p_{24}; and Level 3, System 4 will try to keep p_{24} equal to a constant.

By changing the settings of the Level 3 reference signals, it is possible to change the behavior of the entire hierarchy of control systems. Since the Level 3 reference signals only make sense as 1 or −1, the highest level goals of the system can be changed by entering a new set of Level 3 references, with 1 and −1 assigned to different Level 3 systems.

The perceptual goals of the Level 3 systems (as specified by the reference signals to the Level 3 systems) must be achieved in the external world, as perceived at lower levels of the hierarchy. The state of the external world is determined by the value of the environmental disturbances. Any set of four numbers can be entered as the disturbances. In Figure 5, the disturbances are set to −50, 50, −50, and 50.

Once the disturbance values are set, the model can be run by starting the recalculation process, which is initiated in Lotus by pressing function key F9. Recalculation iteratively calculates values for all the equations in the model; each iteration can be thought of as a change in the state of the entire hierarchy that occurs every fraction of a second.

The recalculation process mimics the dynamic behavior of the control hierarchy. If the model is set up correctly, each iteration of recalculation will bring the control hierarchy closer to satisfying the goal of getting all perceptual signals matching all reference signals. Reaching this goal can take a large number of iterations, especially if you are starting the model "from scratch," so that there are initially very large error signals to be reduced. If each iteration is thought of as representing a brief slice of time (say, .01 sec), then the error signals are reduced rather quickly.

Temporal Resolution

By setting the number of iterations executed per recalculation, you can determine how "fine grained" a look you get at the dynamic behavior of the hierarchy. For example, if you set the number of iterations to 1 per recalculation, then, each time you press F9, you will see the error reduction resulting from each iteration of the control process. If you set the number of iterations to 50, then you will see the error reduction that results after every 50 iterations of the control process. A small number of iterations per calculation is probably best when you are first running the model, because it lets you get a good look at the "dynamic" changes in perceptions, references, and outputs in the entire hierarchy.

Each press of the recalculation key (F9) produces a period of behavior aimed at bringing perceptual signals into correspondence with reference signals. Ultimately, this is done by varying the Level 1 outputs. When running the model, notice that the lower order systems bring their perceptual signals into correspondence with their reference signals almost immediately. Notice also how the Level 2 systems control their perceptions by varying the reference signals going to the Level 1 systems. As the run progresses, the average error at each level of control will oscillate between high and low values. But the system should eventually stabilize, reaching a steady state with very little average error at each level of the control hierarchy.

Disturbance Resistance

When the system reaches a steady state, all errors are at the minimum that can be produced given the current state of environment (defined by the values of the four disturbances). The system has generated outputs that, when combined with the prevailing disturbances, produce values for the input variables that satisfy the reference specifications for perceptions in all systems at all levels of the control hierarchy. The control hierarchy can continue to match these perceptions with reference signals even as disturbances vary. In the real world, disturbances vary constantly, due to changes in the system relative to its environment (as when the system moves relative to a light source) or due to independent changes in the environment (as when a light dims).

The control hierarchy keeps all perceptions at their reference levels by varying its outputs appropriately. You can watch the system solve this problem by typing in new values for the disturbances once the system has reached a fairly steady state. The system will quickly (after a few recalculations) alter all outputs as is necessary to bring perceptions back into line with reference signals. This exercise shows that a hierarchy of control systems controls its perceptual inputs, not its outputs. The reference signals cause the perceptual signals (via the closed control loops) to take on particular values. Thus, reference signals are specifications for input, not output. Outputs depend mainly on environmental disturbances, although some of the output is used to move input variables to new values when there is a change in the reference signal.

The model shows why it would be easy for an outside observer to interpret the behavior of a control system in stimulus-response or input-output terms. Disturbances are the only inputs that can be varied by an agent external to the model. Changes to the disturbances lead to dramatic changes in the output of the model. Indeed, the changes in output are far more noticeable than changes in input, which are negligible. Even though these stimulus-response relationships are dramatic and interesting, they exist because the system is trying to keep various input variables in reference states.

Controlling Behavior

The control system model can be used to show how an outside observer can control the behavior of a control system by manipulating disturbances. The cell labeled "behavior" in Figure 5 contains a number that is a function of the outputs of the system. The number represents a behavioral variable (such as a barpress or rating response) that could be treated as a dependent variable in a standard psychology experiment. These behavioral variables are a function of many individual system outputs (such as muscle tensions and limb movements). You can do experiments to see how disturbances affect this behavioral variable. Just enter values for the disturbances, press F9 to have the system calculate a response, and see the resulting behavior.

The disturbances can be thought of as stimuli that are manipulated by the experimenter. They are the independent variables in your experiments. Once you have a reasonable idea of how each of the four disturbances affects behavior, you can control the behavior (that is, bring it to some desired value) by selecting the appropriate disturbances. The model shows that this approach to the control of behavior works as long the system doesn't change its higher order goals. Try changing the Level 3 reference signal values and see what happens to the system's behavior. This change in behavior would be interpreted as "random error" in standard psychological research. The randomness, however, is only apparent. It results from failure to notice that the system is controlling perceptual variables relative to varying internal references.

Conflict and Reorganization

Once you see how the control hierarchy works "as is," you can do some experiments to see the effect of changes in the hierarchy on its behavior. One simple experiment can be done rather quickly; just change the values of SLOW and GAIN at each level. See what happens when systems at any level respond too rapidly (larger value of SLOW) or with too much GAIN. Test the effect of making SLOW the same at each level of the hierarchy. See if the GAIN can be the same at all levels. Some of these changes will cause the hierarchy to become unstable; the values of signals and variables will overflow their format limits. If this happens, just read the old version of the spreadsheet back into the workspace.

A particularly dramatic experiment involves changing the way one of the Level 2 systems computes its perceptual signal from Level 1 perceptions. This is done by changing the reference to the row in matrix PW for one of the Level 2 perceptual signal cells. A macro has been written (invoked by pressing ALT R) that automatically changes all four Level 2 perceptual computations by randomly selecting new perceptual weights from the PW matrix. Suddenly, the system perceives the world in a new way. If the new perceptions are relatively independent of one another, there will be no problem and the hierarchy will be able to use the Level 1 systems to control them. If, however, the new perceptions are similar to one another, there will be conflict. The control systems will find it impossible to find a set of references for the Level 1 systems that satisfy the references for all the Level 2 perceptions.

When you tinker with the parameters and perceptions of the hierarchical control model, you are playing the role that Powers (1973) attributes to the reorganization system. The reorganization system monitors the status of the control hierarchy; it observes, for example, whether the average error at each level is increasing or decreasing. The goal of the reorganization system (which is a "meta" control system) is to keep the overall level of error in the control hierarchy small (preferably at zero). If error at any level increases (or is too large in the first place), the reorganization system acts on the control hierarchy by changing control parameters and/or the way systems perceive the world. Reorganization does not try to get the control hierarchy to perceive a particular aspect of the environment; it does not even know what the control systems are perceiving. The reorganization system only knows if there is "too much" or "just the right amount" (zero) of error in each control system. Thus, the reorganizing system acts very much like you do when you are tinkering with the control hierarchy. You may not know what the control systems are trying to perceive, but you do know (from the average error) when things are going wrong.

The reorganization system is the learning system of the control hierarchy. You can make learning part of the spreadsheet model by making the reorganization process automatic. This can be done by building a control system that perceives the average error at some level (say, Level 2) of the hierarchy and compares this perception to a reference signal value fixed at zero. The difference is converted into an output which affects the value of GAIN or SLOW of the systems at the appropriate level (Level 2, in this case) of the hierarchy. Once you get this reorganization system to work, you will have an adaptive hierarchical control system—a system that can change its own parameters if it finds itself living in a particularly difficult environment. The development of an automatic reorganization system will not be trivial. An attempt to build an efficient reorganization system for the model is currently in progress.

The Test for Controlled Variables

If organisms are organized as a hierarchy of control systems—and there is considerable evidence that this is the case (Albus, 1981; Marken & Powers, 1989; Miller, Galanter, & Pribram, 1960)—then we can understand their behavior only by learning what variables they are controlling. The spreadsheet model shows that this is not an easy task. Even in a simple environment consisting of only four environmental variables (the disturbances), the system is controlling many perceptual aspects of this environment. An observer can see the system affect many different aspects of its environment. You can see, for example, that the system is affecting the values of the input variables (which are in the environment of the system and the observer). But it is not easy to see that the system is controlling various linear combinations of the input variables. Nor is it easy to see that the system is controlling relationships between different linear combinations of the inputs.

There is a method for determining whether or not the system is controlling a particular function of a set of environmental variables. It is called the test for the controlled variable (Powers, 1979b). A version of this test can be carried out on the spreadsheet control model. You can test whether or not the system is controlling a particular linear combination of the input variables. At the bottom of the spreadsheet model in Figure 5 is a cell labeled "test variable." The number in this cell indicates the row of the PW matrix (row 12 in this case) that will be used as the hypothesized linear weighting of the input variables that is being controlled. The four weights associated with this test variable are shown in the row of cells labeled "weights." To test whether or not this variable is controlled, we produce a sequence of disturbances to the four input variables. If the test variable is controlled, the disturbances will have far less of an effect on this variable than would be expected.

The expected effect of the disturbances is measured in terms of the expected variance of the test variable. If the test variable is not controlled, its expected variance is equal to the sum of the variances of the disturbances and system outputs. If the test variable is controlled, its actual variance will be considerably less than expected. The ratio of expected to observed variance is called the stability factor. It is a measure of the system's ability to stabilize (control) an environmental variable against disturbance. The stability factor is shown at the bottom of the spreadsheet model (Figure 5). Its value will be close to 1 (expected equals actual variance) if the test variable is not controlled; the test variable is not stabilized. Its value will be close to zero if the test variable is actively destabilized; the system increases the variance of the test variable. The larger the value of the stability factor, the more likely it is that the hypothesized controlled variable is, indeed, controlled (stabilized against disturbance). In the present case, even test variables that are not controlled will have a stability factor that is considerably greater than 1.0 (often close to 150.00), because these test variables are very similar to the controlled variables. To do the test properly, several different variables should be tested. A controlled variable will have a stability factor that is two to three times larger than the stability factor for an uncontrolled variable.

The stability factor is computed for a sequence of 15 sets of disturbances by running a macro that is initiated by pressing ALT S. After each set of disturbances is applied, the model goes through several iterations (the default is 50) in order to bring the input variables to their reference states. The actual variance of the test variable used in the computation of the stability factor is based on the steady state values of the input variables. After a transient disturbance, the control system requires several iterations of calculation to reach a steady state.

DISCUSSION

A method of implementing a hierarchical control system model on an electronic spreadsheet has been described in this paper. The spreadsheet model makes it possible to explore the behavior of a hierarchy of control systems in some detail. A number of experiments are described that can be done to see how the model responds to changes in the environment and in its own higher level reference signals. The spreadsheet implementation makes it relatively easy to change various aspects of the model, such as the control parameters and perceptual functions. It is also possible to expand the model, by adding levels of control or by adding systems at each level.

Application to an Experiment

The spreadsheet model described above simulates the behavior of an organism that is controlling four intensities (Level 1), four sensations that are linear combinations of these intensities (Level 2), and four relationships (Level 3). It is not a model of behavior in a particular experiment. But it is possible to augment the model to simulate behavior in an experiment. For example, a version of the spreadsheet model, with only two levels and two systems at each level, can produce coordinated behavior like that observed in experiments described by Marken (1986) and Bourbon (1989). The input and output values of the spreadsheet model can be compared to the values observed when subjects control variables such as the position of and distance between lines on a video display screen.

Spreadsheet Modeling

The electronic spreadsheet provides an excellent environment for quickly implementing and displaying the results of working models of behavior. It was possible to write and debug the three-level hierarchical control model in less than an hour, a process that took three times longer with a more conventional programming environment (PASCAL with a fast compiler). Moreover, once the spreadsheet version of the model was completed, it was easy to change (for example, by adding systems or levels). It is much more difficult to make such changes in the conventional language version of the program.

One of the main advantages of the spreadsheet approach to modeling is the ready-made display format. When developing the control hierarchy model, you do not need to write code to display model values at appropriate locations on the screen. One of the main disadvantages of the spreadsheet approach is that the code is run interpretively; the speed of execution of the spreadsheet model is fairly slow. This would be a significant disadvantage if the goal of the control system model were to simulate behavior in the context of continuously changing environmental disturbances, such as those used in many manual tracking experiments. Simulation of behavior in these circumstances is best left to more powerful, compiled versions of the control hierarchy.

REFERENCES

ALBUS, J. (1981). *Brains, behavior and robotics*. Petersborough, NH Byte Books.

BOURBON, W. T. (1989). A control-theory analysis of interference during social tracking. In W. Hershberger (Ed.), *Volitional action: Conation and control*. Amsterdam: North-Holland.

FOWLER, C., & TURVEY, M. (1978). Skill acquisition: An event approach with special reference to searching for the optimum of a function of several variables. In G. Stelmach (Ed.), *Information processing in motor control and learning* (pp. 1-39). New York: Academic Press

LEBLOND, G. T., & COBB, D. F. (1985). *Using 1·2·3*. Indianapolis IN: Que Corporation.

MARKEN, R. S. (1986). Perceptual organization of behavior: A hierarchical control model of coordinated action. *Journal of Experimental Psychology: Human Perception & Performance*, **12**, 67-76.

MARKEN, R. S. (1989). Behavior in the first degree. In W. Hershberger (Ed.), *Volitional action: Conation and control*. Amsterdam: North-Holland.

MARKEN, R. S., & POWERS, W. T. (1989). Levels of intention in be

havior. In W. Hershberger (Ed.), *Volitional action: Conation and control*. Amsterdam: North-Holland.
Miller, G. A., Galanter, E., & Pribram, K. H. (1960). *Plans and the structure of behavior*. New York: Holt, Rinehart & Winston.
Powers, W. T. (1973). *Behavior: The control of perception*. Chicago: Aldine.
Powers, W. T. (1979a, August). The nature of robots: Part 3. A closer look at human behavior. *Byte*, pp. 94-116.
Powers, W. T. (1979b, September). The nature of robots: Part 4. Looking for controlled variables, *Byte*, pp. 96-112.
Powers, W. T. (1989). *Living control systems*. Gravel Switch, KY: Control Systems Group.

NOTES

1. For clarity, spreadsheet cells are referred to by functional names (e.g., O and R for "output variable" and "reference signal" cells, respectively) rather than spreadsheet coordinates (e.g., A15 and A12).
2. The model is available from the author as a Lotus 1·2·3 worksheet. The worksheet will be sent upon receipt of a formatted 5¼- or 3½-in. double density or high density disk in a reusable mailer with return postage.

(Manuscript received April 13, 1990;
revision accepted for publication July 16, 1990.)