Volume 1

Number 4

Published by the Control Systems Group

Closed Loop # 4

Threads from CSGNet

CSGNet, the electronic mail network for individuals with control-theory interests, is a lively forum for sharing ideas, asking questions, and learning more about control theory, its implications, and its problems. The following "thread," stitched together from just one of the Net's many ongoing conversations exemplifies the rich interchanges among Netters.

There are no sign-up or connect time charges for participation on CSGNet. The Bitnet address is "CSG-L@UIUCVMD" (use no quotes in this and the following addresses); "CSG-L@VMD.CSO.UIUC.EDU" is the Internet address. Messages sent to CSGNet via these addresses are forwarded automatically to all participants. Via CompuServe, use the address ">INTERNET: CSG-L@VMD.CSO.UIUC.VMD" to reach the Net. Initially, you should send a note to the network manager, Gary Cziko, at "G-CZIKO@UIUC.EDU" (Internet) or at "CZIKO@UIUCVMD" (Bitnet); Gary's voice phone number is 217-333-4382.

Greg Williams
606-332-7606
October 1991

A (Control) Engineer among the Psychologists

Izhak Bar-Kana: My main field is adaptive control, with some neural networks and robotics. I subscribed to CSGNet mainly due to its "control" name and intended to be a quiet listener. The discussion is very interesting, and to me, quite surprising.

Regarding positive feedback and evolution, isn't "adaptation" the word? A simple time-invariant mechanism can perform that much. If the task becomes more difficult, "tracking errors" (based on the performance index the system tries to minimize) are used to change the gain, for example speed vs. error (as in the case of a pursuer), even in the simplest adaptive mechanisms. After a while, what previously was an extremal situation becomes a normal situation, because the evader is faster. A learning system identifies it as the normal situation, which in my humble opinion is expressed by development of muscles, etc., as when we train. But this is only a pretext to introduce myself, and to explain why I will be mainly a quiet listener, at least for a while, until I get the opportunity to read more of your works and speak, at least, the same language.

Wayne Hershberger: Welcome, Izhak. It sounds to me as if you already talk the CSG "language" very well, if I understand you correctly. I would put it this way: adaptation, including evolution, does not require positive feedback; rather, positive feedback requires (calls for) adaptation! What adaptation requires is *random* (polarity) feedback. That is, adaptation can be viewed as "slow" control (long-loop time) in which elements in the system periodically render the polarity of the feedback loop *random,* in order to "discover" the arrangement of subordinate elements which will restore the level of the system's error signal to within tolerable limits. An *E. coli's* control of its locomotion is the canonical (or prototypical) case of such slow control, complete with a nearly literal, random, "roll of the dice." By varying the rate at which the *E. coli* tumbles (or the rate at which Darwin's blind variations or Thorndike's trials and errors occur), the polarity of the system's feedback can be either (a) maintained when it is negative, or (b) changed when it is positive (and, thus, eventually restored to negative). The adaptation of organisms to overwhelming disturbances is the restoration of control. In his *Design for a Brain,* Ashby called it ultrastability. Bill Powers has called it reorganization. Perhaps it should be called "slow control utilizing random-polarity elements which ensure that feedback is negative in the long run."

Whether evolution (the adaptation of species) is "slow control" is a matter I will leave for others to decide. But the adaptation of organisms is clearly "slow control."

Tom Bourbon: Izhak, welcome, even if you hope to remain a listener. Perhaps we can coax you from that intention. Many of us are in the behavioral and social sciences, so we lack backgrounds in your areas of expertise. I, and probably several others, would appreciate information from you about good general references on the topic of adaptive control. I am presently working on models of human tracking behavior in which two people, or a person and a control-system model, interact. Two people can easily decide to change from one mode of interaction to another, and one person can easily recognize when the other has changed, then adapt to the new mode. I want my modeled person to develop the same capacity as a real one who detects the mode employed by the real person, then adapts. I'll admit I'm in over my head on the topic of adaptive control, but I suspect there might be some basic ideas there that will help me in my work.

I hope you will reconsider your decision to remain silent—yours is precisely the kind of expertise many of us lack!

Izhak Bar-Kana: Many thanks to Wayne Hershberger and Tom Bourbon for the welcome. It is not easy to keep quiet in such an active environment, though I think I must do a lot of reading and listening to you before I even understand you.

To Bill Powers (and actually to all): I am asking more than claiming, but I am not sure I can agree with the apparent contradiction between engineering control diagrams and living control systems. Or, better, I do not understand it. If the problem is driving a car, the input is the way, the trajectory which must be maintained, and the output is the position of the car. Of course, this difference must be measured, and the control system only receives the output of the sensor which measures this difference. In ideal situations, this measure is exact. In other conditions it has noise, bias, miscalibrations, phase lags ("time constants"), and/or transport lags ("pure delays"). The control system tries to bring the error signal to zero, and the output is the position of the effector ("actuator"). Between the sensor and the effector (motor) there is a controller which transforms the signal in such a way that stable performance of the control system is guaranteed. And this is only the simplest control system. If a "brain" is involved, the signal transmitted to the effectors cart take more sophisticated forms: the brain might know the performance of the control system, might be capable of taking into account its time lags, etc. Furthermore, the brain has stored the final aim of the trip, and might change the route or make other decisions which could not be taken by a simple autopilot whose only purpose is to keep "in line." But I think there must be some separation of the various tasks. And even here, the final point is stored in the brain only because some real final point is in the real world, and this is what we call "input," even though the control system can only affect the output of its own sensors, or its perception of the real world. If the temperature must be maintained, the input can be considered internal, because it starts in the brain. Still, this signal is transmitted to a control system whose function is to execute and reach this temperature, or to annihilate the difference between the desired temperature (registered in the brain?) and the temperature of the body. For this control system (or better, regulation system) the input is external.

I don't understand how the living control system affects its inputs. Maybe only a difference in definition? In a tracking system, the position of a target is the input; the resulting position, of the eye for example, is the output, even if the only physical and measured signal is the difference between these two values. I agree with everything I can claim I understand in Powers' "Manifesto," so maybe I miss the main point here. I would appreciate if you could open my eyes here, because

I am trying to understand, not to prove that I am right.

To Rick Marken: Maybe the engineering control people need other tools because they must *design* the control systems, not only understand them. The "sophisticated" control people use lots of math because of the difficult task of proving that a system is stable. Not because they are crazy about stability, but because it is easy to get an unstable system with an "ingenious" and "intuitive" control method. When control is nonstationary and nonlinear, such as in adaptive systems, the problems and the proofs are even more difficult. The problem is that if you cannot know (prove) that an adaptive system (I mean "engineering" adaptive mechanism) is stable, in general you will discover that it is unstable under some conditions. I don't know how much this group is interested in or how much time it has to spend on this stuff, unless people are interested in the instability mechanisms of pathological cases.

Please see my lines on the car driver above. Of course the organism only receives the signal supplied by the sensors, but that is more or less the measure of the external signal. By the way, besides dealing with theories of systems, I am also an engineer, and I can tell you that no engineer would let a motor run, much less a plane fly, without thousands of simulations, no matter what the theory says, and in fact the theory, the complex functions, differential equations, etc., do not say much when a real, large, complex system is involved. And I would not dare to compare any complex plane with a living organism, not to mention an intelligent creature. So, learning through observation and simulation is a main engineering tool. But when I want to *design* a stable and well-behaved system, I need mathematical tools which express stability and performance, and their dependence on the various parameters I might or might not change. And then things start getting tough, like trying to define pornography: It is hard to define, but it is easy to recognize when you see it. Yet I usually need the differential equations to have reliable simulations, especially if I want to discover when the real plant stops performing satisfactorily. It is not that important whether your simulations are state-of-the-art or not, as long as they are correct and approximate the real thing. I don't know your models, so I hope they are.

To Tom Bourbon: At this stage, I am afraid I can only tell you that the problem is interesting, and that I only have begun studying it. It is not as much an adaptation problem as it is a learning problem. How to guarantee that a mechanism learns while it performs its task and maintains a stable behavior is not an easy task! I will try to be more specific in the future. In fact, part of the new trend in "intelligent (automatic) control" tries to eliminate the differential equations because "the brain does not solve differential equations," and tries to imitate the brain; the algorithms used are just (poor) attempts to reproduce the activities of organisms' neural networks.

Bill Powers: Izhak, I think that the mental model you are using is the one traditionally given in engineering texts, the same one that Norbert Wiener picked up and used in his first book on cybernetics. In that model, "input" means *reference* input. It is shown, usually, entering the comparator as if from the external world. The feedback signal, on the other hand, is just "picked off" the output variable through some feedback transducer.

In the model we use in the Control Systems Group, when we say "input," we mean the sensory feedback input, not the reference signal. That is because the sensory inputs in the organism constitute the "feedback pickoff" which reports, as analogue signals, the states of external controlled variables. The senses do not report the intended or desired state of affairs; only the current actual state of affairs. The reference signal comes not from outside, but from systems superordinate to the control system in question, inside the organism.

So it comes down to how we match the main functions and signals in a generic control system to corresponding functions and signals in a particular control subsystem in the organism. The abstract organization is the same; our model has the same connectivity as the one I believe you are using, so the control-system analysis itself is unchanged. But the meanings and the implications are greatly changed.

Apply this to a model of driving a car. The driver sees the current position of the car relative (laterally) to the road. Out of all of the information in this image, the brain extracts a position signal which varies as the car moves from side to side. Thus the position of the car is the input variable, not the output variable, in the steering control system. The position signal is compared with another signal which specifies the intended state of the position signal: that, of course, is the reference signal. The driver can select any possible perceivable position on (or off) the road as the reference position. The error signal, reference minus sensory signal, actuates the output of the control system, which is the torque applied by the arm muscles to the steering wheel (this requires two phase advances for stability). That torque is the last variable in the output chain which is due strictly to activities in the brain. From there on, we have mechanical linkages and external disturbances coming into play, which alter and add to the effects of the output and are not themselves part of the behaving system. The result is some position of the car on the road, and thus a state of the perceptual signal representing that position. The feedback effect keeps the perceptual signal in a match with the reference signal (give or take dynamic and static errors). It is not necessary for the brain to contain any detailed knowledge of physical properties and events outside itself other than the controlled variable. Variations in output properties have little effect; disturbances are automatically counteracted without any need to anticipate or sense them (except through their effects on the perceptual signal).

With the reference signal moved inside the control system, we can now "parse" complex behaviors in a new way. In order to alter the position of the car on the road, the brain now needs only to alter the reference signal for the steering system that is now in place. To pass a car, higher systems concerned with relationships to other objects change the reference signal enough to move the car to the other lane, keep it there a while, then move it back. Of course there is also a speed control system operating independently, with its own input which senses speed and its own actuator which affects speed (the foot on the accelerator pedal). The "passing-another-car" system alters the reference signal for speed, too, as the driver passes the other car, first increasing it, then decreasing it. So the higher system uses, the lower systems by manipulating their reference signals, just as a human user manipulates an artificial control system by turning the knob which changes its set point.

At the same time this is going on, the driver can use one arm and hand to reach out and change the volume on the car radio, then scratch his neck, all while telling a joke to the passenger. In this model there are many control systems acting concurrently, each controlling just one (perhaps complicated) variable. There are neuroanatomical justifications for breaking down behavior into multiple control systems operating independently and in parallel, and organized into levels of control. And I think this picture also helps us to approach the modeling of complex behavior in an orderly way, solving problems of peripheral control to serve as the foundation for exploring systems at higher levels, more central in the nervous system. We can, of course, pick isolated systems at any level and analyze them as control systems, absorbing lower-level control system properties into their output functions. But the final model must spell out all of the stages of control which exist, while, one hopes, maintaining correspondence to known structures in the nervous system.

The same model applies to human temperature control. The input variable is the temperature of a sensory ending (in the hypothalamus, I think). The reference signal is variable, as temperature can be maintained actively anywhere between 98°F and about 104 or 105°F. I don't know what varies the reference signal, although I know it changes when you get sick. The error signal is translated into shivering and peripheral vasoconstriction if it is positive (sensed temperature lower than reference temperature) and into sweating and peripheral vasodilation if negative: that is the behavior which affects the input, the sensed temperature.

As for simulations: we use them a great deal, where we know how to construct them. They work very well. Stabilization has not yet proven to be a problem, although in the arm model you have seen mentioned, the problem was solved just by introducing known properties of the neuromuscular systems in question (we never set foot on a complex plane). I think that the hierarchical structure simplifies stabilization problems, which might be an indication (and might not) of why the whole system is organized that way. We haven't gotten very far with modeling very complex or high-level behaviors. We're still taking baby steps and learning how to walk. But I think that our approach, probably combined with some of the perceptual models being developed by neural network people, will carry us a good deal further before we have to change the basic structure of the model.

Izhak Bar-Kana: If this is a family with specific definitions and problems, which I do not belong to, I would rather keep quiet or say "excuse me, it was nice meeting you." But if it is a control-theory group, and when the thermostat is an illustration, then some things deserve, at least, clarification. The model you all seem to use, position = position + k x error, is of course, perfectly correct. However, I am confused about what you call output, controlled object, and control objective.

No matter how good or how bad is the function of the organism, or of any control system, the only thing which it can affect is the output. Even if one closes one's eyes, or if one is drunk, he or she still controls the output, which is the name of the controlled variable in my dictionary. A closed-loop system, properly designed or properly organized, will try, in general, to minimize the signal representing the error between the measured reference signal and the measured output (controlled variable). But this is not the only way to do the job.

A fresh driver, on a new route, with a new car, performs very much as was said above. However, a better controller uses all possible prior knowledge to get a "better" quality of control. In our case, the effect of learning (and here we very much try to learn from organisms) seems to become cardinal. After a while, the brain has a sufficiently good model of both the route and, for example, the car, and the control is very much open-loop. Based on some details of the route, the brain predicts both the changes and the rate-of-change of the route, and also the response of the car to various inputs, and passes to the car a very complex signal, mainly open-loop control. No, I do not try to advertise open-loop control! There is still much uncertainty in my model of the route and in my model of the behavior of the car, therefore I keep my eyes open and monitor the error between the desired trajectory and the actual trajectory, but now the gain of this error can be much lower, and my mainly open-loop control is now smoother and "better" than the closed-loop only. In both cases, we control the position of the car. The input is the controlling, not the controlled signal, even though in closed-loop it might be hard to tell.

Now, about the external reference. Indeed, no route can tell me where I want to go, which is a decision. Once I decide where I want to go (and, hopefully, based on the knowledge or valid "representation" of what is going on out there), I must follow a route which exists somewhere (whether I follow it or not). This is the reference input which is measured and transmitted, after "cleaning" (filtering) it of noise and processing it so it fits the needs of the control system (the transfer function with its leads and lags, nonlinearities, etc.). At the same time, the "result" (as named in Bill's answer to me—why not "output"?) of the control, the position of the car on the route, which I call output, is measured and compared with the reference. The difference is then the controlling or one of the controlling signals which now affects the output. The objective of the control can be to cancel this error, or to minimize it without requiring the use of all vital reserves, or to minimize some combination of the error and other variables.

What is wrong here? I don't use this model because it is Wiener's model, or anyone else's model, nor because it fits some complex mathematical formulas, but because that is how I understand the control systems. If a thermostat is designed to maintain a constant temperature, then it is a regulator, and the referenced temperature can be considered as being internal, corresponding, I hope, to some desired but real temperature. Good or bad, the system can only control the output, the heat in the room; either it measures it correctly or not.

Sometimes here I seem to get old names for old things, only shifted. For example, Bill, what does it mean to say that "position is represented as integral of motion?" Position is the integral of velocity. I am not surprised that in the brain, motion is hierarchically higher than position, because motion is obtained by changing positions.

Bill Powers: Izhak, we do have some work ahead if we're to achieve communication. I'm sure that when it comes to control theory itself, you are the expert and I am the amateur. But I don't think we will have any arguments in that area. Where we differ is in how we analyze the behaving system itself, prior to setting down any equations. This leads to some differences in terminology, which we can certainly straighten out. But there are also some fundamental assumptions which do make a difference, and we might have some long discussions about those.

Let's start with the thermostat, honoring tradition. I partition any control system into the system itself and an environment. Where you draw the boundary is arbitrary, as you must have a closed loop in the end, but I think it's worthwhile to give some thought to the natural boundaries. In the thermostat, the bimetallic element is the system's sensor of temperature, so that is a natural input boundary. The contacts actuate a relay, which turns on the furnace, which converts a supply of oil or gas into a thermal output; that is another natural boundary. I call these "natural" boundaries because they separate signals and functions in the controlling system from processes in the independent environment. The sensory signal (position of the metallic strip) depends only on the temperature of the sensor itself. The thermal output of the furnace depends only on the stored energy and the operation of the furnace. All other processes between the thermal output and the sensor (in the external world) are subject to potential disturbances and changes of parameters, and those changes occur independently of what happens "inside" the system, as I have defined it.

Now, with these definitions, what is it which the thermostat controls? That means, what variable is held most nearly at a specified reference level, in spite of all kinds of changes which can occur in the independent environment? Clearly (to me), that is not the temperature in the rooms of the house. If a bedroom window is open, the bedroom will become cooler. If a window is open in the living-room where the thermostat sensor is, the air near the window will become cooler, but the air immediately around the sensor will be maintained at the reference temperature (plus or minus the dead zone). So we know that the temperature of the air immediately around the sensor is closer to being the true controlled variable than the temperature of the air on the other side of room or in a different room of the house.

Now suppose we set an infrared heater on the floor, so that its radiation passes through the grille of the temperature sensor and falls directly on the bimetallic strip. What will happen? The contacts will remain open and the furnace will remain off until the temperature of the sensor element has dropped once again to the reference temperature. This will happen when the air around the sensor has cooled enough to remove heat at the same rate that the infrared radiation is adding heat. So the room temperature will drop and the temperature of the air at the sensor will drop. What will remain the same? *The signal (position) which stands for temperature.* The temperature of the sensor itself will be held near the reference temperature—but nothing more remote from the sensor will be controlled. By varying the infrared heater's output, you can cause any air temperature you like, and the control system will still experience zero error.

So by thinking of various ways to disturb the temperature outside the system, we show that the only variable reliably stabilized against independent perturbations is the state of the sensing element itself, and of course the signal associated with that state. Variables more remote from the input boundary of the system will be stabilized by the same control action only if they are shielded from extraneous disturbances by something other than the control system, and only if the sensor's temperature depends reliably on those remote variables.

I think that if you analyze any control system carefully, you will see that the same principle applies. The only variable stabilized against independent and unpredictable disturbances *of all kinds* in the environment is the input signal of the control system.

Now, the output. You suggest that the output of the thermostat System is "the heat in the room." But if that is considered the output, then we are in the position of not being able to predict the thermal output of the control system's furnace, because the heat in the room depends on many factors other than what the furnace is doing. The outside temperature, sunlight coming in the windows and falling on outside walls, open windows, fires in a fireplace or a stove, lamps, and blockages of air circulation all contributed to the heat in the room. The furnace's thermal output also contributes to the heat in the room, but it does not *determine* the heat in the room. The only thing which the control system can *determine* is the amount of heat released in the furnace by burning fuel. So I would call the output of the thermostat the quantity of heat released (per hour) by combustion in the furnace. What happens to that heat after it is generated, what other sources and losses of heat might be acting at the same time, is impossible to predict from knowing how the thermostat is designed. The only "output" which depends *strictly on the operation of the control system* is the thermal output of the furnace.

Between the thermal output and the temperature of the sensor, we have a variable and unpredictable environment. The control system is not equipped to sense any of the causes of those variations. Nor does it need to be so equipped. All it needs to "know" is the temperature of its own sensor. By knowing that, and by being able to vary its contribution to the temperature of its own sensor, it can control that temperature.

This is the basis of my general approach to analyzing control systems prior to reducing the relationships to equations or programs. What do you think of it so far?

Izhak Bar-Kana: Bill, thank you for the compliment, but you know and I know that one must learn a lot to discover that no one is expert. In fact, I try to understand your discussions as control theorists in psychology because they might be very relevant to my questions in artificial learning. I am not sure I really understand the difference between my model and yours; maybe what you call "input," I call "measured output."

By the way, the temperature sensor is supposed to be in distant rooms, not in the neighborhood of the furnace. Now, if the window is open near the sensor, the temperature in this room is not affected by the thermostat system. In this case, the system receives some value of a low temperature, sends out lots of heat, the only thing it can directly control, and has no effect on what it is supposed to affect, the desired output, the temperature in the room, or on its measurement, the feedback input to the system. It does not change the fact that the input signal is used to affect the output or different stages of outputs, some of them measured if needed, to guarantee satisfactory behavior of the system. The basic design would take into account the basic thermal properties of the room, and nominal ambient external and internal temperatures, and the rate of fuel burned and heat supplied would have to maintain this nominal condition "almost" with no other regulation. So, the input is the desired temperature, and the output is the actual temperature. Because I know that uncertainties are always present, I monitor the actual "controlled output" and use the difference between the desired output (by the way, I also call it the reference input, so maybe I just live in another kind of dichotomy) and the real output, and use it (after filtering the measurement noise) to generate supplementary heat that, hopefully, will compensate not only for uncertainties, but also for changes in the ambient temperature, and so on. I might even decide that the closed loop is enough, especially if the desired temperature is fixed. The error between the desired and the measured output is used as input to the controller, amplified and processed, and then sent to fix or change the rate of fuel, the rate of heat, and change the temperature in the room. If the loop is well-designed, it finally brings the room temperature (as it measures it) to, or close to, the desired temperature.

I might repeat myself, but now suppose that the sensor is broken and frozen at a fixed low temperature. Then the control system gets some constant input and sends waves of heat, changing, I call it controlling, this output, whether this is the desired output or not. Anyway, I think I start to understand you, and I am only afraid that it might be difficult having a dialog with the control community at large, if the claim "the control system controls its input" is not understood. In spite of this argument, I think that things are not as distant as they might seem. Some more eavesdropping and more reading from my side will straighten out things even more. I am not used to such a high level of patient discussion about different opinions, and I am honored to participate in it, even if I might introduce noise.

Bill Powers: Izhak, I am very pleased that you are so open-minded and willing to work out these problems of communication. As we go on, I'll try to address selected topics in your communications which might help resolve whatever problems there are. Today's text: "I am not sure I really understand the difference between my model and yours;

maybe what you call 'input; I call 'measured output."' I think you have the key to one of our differences in nomenclature. What I call input is what you call measured output. I call this "output" an input partly because it is the external variable which affects the sensor, an input to the control system, and mostly because it is affected only indirectly by the actual output of the system—that is, the actuator or effector. The state of the controlled variable is not determined by the system's action alone; control is required because there are unpredictable disturbances which also contribute. If we use the term "output" for the controlled variable, then we have the odd circumstance in which we can't define the output of the control system itself—the output depends on independent factors just as much as on the behavior of the effector. I prefer to reserve the term "output" for the effector's action, which is the last thing in the chain of output processes which is completely determined by the control system. Between the effector's output and the controlled variable there are many sources of disturbance—they are the primary reasons control is needed in the first place. We can agree on the term "controlled variable." But I claim that variable is more closely associated with the sensory input than with the effector output.

I don't think it would be practical to design a home thermostat as a basic open-loop system with feedback added to handle details. The "details" are the whole problem. The steady-state thermal output of the furnace which is needed ranges from zero to the most the furnace can produce, depending on unpredictable heat losses and gains in various seasons and at various times of day and night, varying conditions of cloud cover and wind velocity, and various conditions of occupancy. A real home thermostat is simply designed by picking a furnace which can keep the room above the maximum desired temperature on the coldest cloudy day at 100 per cent duty cycle, and then letting the feedback do the rest, as it does. The problem a thermostat has to deal with isn't "uncertainty" in the sense of system noise. It's the fact that there are very large and unpredictable disturbances of the controlled variable. When the main causes of variations in the controlled variable are major disturbances rather than noise, there's no way an open-loop branch in the system can accomplish much.

Human systems, I believe, are in the same situation. Most of the output (motor) activities which take place are there to counteract large disturbances of controlled variables. There is very little random variability in the system itself—only a few percent of the range of action. Living systems got the reputation of being highly variable because the wrong model was applied: psychologists thought that the behaviors are responses to stimuli, whereas they are probably just actions which protect controlled variables from disturbance. The disturbances, of course, were mistaken for stimuli. Because disturbances occur unpredictably, the behavior counteracting their effects is just as unpredictable. But if you know what variable is being controlled, behavior suddenly looks far more regular: it opposes the effects of disturbances in a highly systematic way.

Suppose you have a motor controlling the angular position of a load through a gear train. The angular-position sensor is located at the load and not on the driving shaft, because the shaft can twist and there can be play and runout in the gears. The actual output of the system is not the position of the load, but a torque applied to the armature. If an extraneous force is applied to the load, the torque will immediately rise to counteract it. If the force is too large, the motor still produces maximum output, but the shaft does not turn: the position no longer is affected by the system's output torque. If we now remove the position sensor and substitute a tachometer, without changing anything else (except perhaps the stabilization filter and scaling amplifier), the controlled variable becomes angular velocity instead of angular position. The output torque, in the steady state, equals the sum of all frictional and viscous resistance plus any opposing torques, and the angular velocity matches the reference signal. So *the nature of the sensor determines the nature of the controlled variable.* That's another reason for saying that the controlled variable is an *input* variable.

The reference signal is certainly one input to the system's comparator. The sensor signal is another. But I like to say "reference signal" because in living control systems, reference signals very often are supplied by higher-level systems, not by sensory inputs. In fact, I can't think of any case where a known reference signal comes from the sensory inputs. But we can get into that later. Is my nomenclature beginning to make any more sense to you?

When an engineer designs a control system, there is a control problem defined in advance. The engineer can see what effectors are needed, what sensors are needed, and even (sometimes) how control might be achieved by open-loop means (if there aren't any unpredictable disturbances of the controlled variable). That's because the engineer has an internal reference defining the desired result and the ability to shape a device having known effects on the physical world. The engineer can see all the inner details of the control system, and he or she also has advanced knowledge of the properties of the physical world with which that system will interact.

An evolved organism like a human being, in the process of becoming organized, doesn't know any of that. The environment is known only through sensory inputs and direct physiological effects of the environment on the body (the state of which is also sensed or at least sampled). Actions affecting the environment are produced only by sending signals into muscle systems. There is no a *priori* information in a growing brain concerning physical properties of the environment or the body, or any 'laws of nature" or general principles. The brain can't use any of the engineers knowledge in building up its own control systems as it matures.

Whatever the brain does, it must do on the basis of available information and whatever amount of organization it has and has acquired at a given stage of development. It discovers properties of the external world relevant to control only by acting and sensing the result. No information is available about what happens between action and sensed result: only the result is experienced. Many consequences of acting are insignificant and unrepeatable; only some are consistent and therefore possible to control. The organism learns about the consistent sensory results, selects remembered states of those results as targets to repeat, and by trial and error finds the combinations of output acts which will tend to restore the sensed world to any former state, starting in any other state. That is called control. It does not need to know, and in most cases never does learn, *why* performing a given act results in a given effect on perception. It does not need to know the true nature of Reality, explaining why perceptions behave as they do. Only when higher-level cognitive systems develop does a brain begin to acquire a symbolic understanding, a model, of the external world, so it can explain why acting in a certain way is necessary if control of a certain kind of experience is to be possible. Only at that stage can someone become a control-system engineer—and by that time, the vast majority of the engineer's own personal control systems have been in place for years.

If you ask a child, or for that matter almost any adult, *why* turning a certain knob on a television set makes the picture get brighter, the answer is going to be something like, "Because that's the brightness control." But nobody has trouble with adjusting the brightness to whatever level seems "right." Control is not based on understanding of the physical world. That's lucky for you, if you are a bacterium or a baby.

Here's my nomenclature for a living control system:

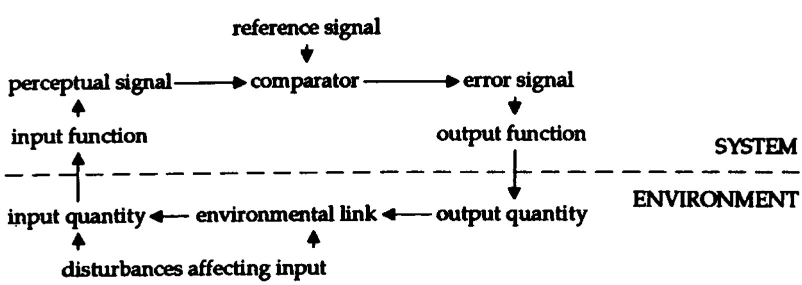

In the diagram, "quantity" = physical variable outside system, "signal" = physical inside system; "output function" = effector, "input function" = sensor + sensory computations, "output quantity" = direct measure of effector action, not subject to external disturbances, "input quantity" = physical variable directly affecting sensor, "disturbances" = independent contributions to state of input quantity, and "environmental link" = path by which effector contributes to state of input quantity. The observable controlled variable is the input quantity. The output quantity varies as disturbances vary, and cancels their effects. Note that the reference signal originates inside the system.

Izhak Bar-Kana: Again, trying to read and understand the discussion which is going on, I have the feeling that I broke into a subject in which I cannot even call myself a novice. I am very interested in your discussion, because I want to use any information related to biological intelligence for our "engineering" intelligent controls. I am pleased and surprised with the nice and patient tone of this discussion, and with the detailed explanations I get.

1. My position in control is actually very similar to yours. I am an engineer, a control one, and mainly, a feedback-control one. In what I call "simple adaptive control," I have lots of fights with some colleagues, *because* I show that some simple feedback loops with "appropriate" adaptive gains can perform much better and can be more dependable than many sophisticated schemes.

2. I cannot afford to mix control with philosophy. For control jobs, I take external "reality" for granted. When my car does not crash into a wall or another car, I deduce that my sensors' approximate interpretation of "reality" was not very bad, which I am not sure of if I am drunk. In both cases I control my sensorial perceptions. I am glad that I am around here and now, even though this phenomenon was not present some time ago and will vanish sometime in the future, and in spite of the fact that I am "mostly void" and from what is left, more than 90% is just plain water. The car which comes toward me might discontinue this experiment too soon, even if it (the car) is also "mostly void."

3. I take for granted that you are there, and that I read a translation of your thoughts, even if my perception might be very poor.

4. As I understand it, the great idea of Bill Powers is the observation or the discovery of *feedback* control in behavior, as opposed to a simple stimulus-response interpretation. I could not agree more.

5. I might make people angry, but because feedback is closed-loop, it is not very important to me what you call what. I only need the dictionary.

6. I do not confuse effect and control. I assume that in order to control, one must be able, first of all, to affect, and in a desired mode, too.

If the thermostat cannot control the temperature in the room, how can it control the measured temperature, which is its input? I hope you do not translate “is affected by” into “controls.” I assume the thermostat system is designed to perform a reasonable job, as our behavioral control loops are “designed” or have developed reasonable tasks. The others are not around to testify.

7. One of my problems is that I am, in general, in the middle, and I try to understand all sides, and get the main points of opposite ideas. I am afraid that this group, trying now to explain every behavior as closed-loop, might ignore some very nice and intelligent open-loop controls, based on the splendid property of the brain: learning. Yes, closed-loop is dependable, and when one learns skating, one uses very stiff closed-loop control. With time, after one learns one’s own behavior and the response of the skates, one uses much open-loop control, based on the learning process and modeling of this behavior. Closed-loop control is still there, but not alone, and not as stiff as before (“lower gains”). A predictive closed-loop control is also there, comparing the predicted desired position to the actual present position, but this loop is weaker and weaker with training. How much is open and how much is closed is only a question of control gains. It is not a cause-effect interpretation. It is very intelligent open-loop control. In my humble opinion, one misses something if one ignores this aspect of control, in particular when learning (I wish I knew how) is involved.

8. About “controls output” or “controls input”: as I said, it is not very important (for me), if I know what we are talking about. Yet, many of your conclusions are based on modeling. Control experience tells me that I cannot derive conclusions about the behavior (or transfer function) of components of a control loop from the behavior of the closed loop, *because* of the nice property of feedback, namely, lack of sensitivity to variations in the parameters of various components. Therefore, to find the transfer function of the *closed* loop, I must first *open* the loop, test the input-output behavior of the *open* loop, when any variation of the components expresses itself one-to-one, and only then close the loop. The math also shows that analysis (at least Gedanken) of an open-loop system provides conclusions relevant to the behavior of the closed loop, but this is only secondary here. Anyway, this might explain why I have an input and an output (or more outputs on the way, but only the controlled variable is the output of interest). Now, I measure the output, and the result of measurement becomes the feedback input, and the other input is the reference input. Along with the closed-loop scheme, which is basically identical to Powers’ scheme, I add a function of the desired (reference) input which directly drives the control system, bypassing the comparator.

Bill Powers: Yes, Izhak, it's hard to convince some people that a simple control system can accomplish more than many extremely complicated approaches to the same task. One of the great difficulties is getting people to think in terms of continuous variables, isn't it? The digital revolution really brainwashed everyone. Even in electronics, technicians are happy to learn digital circuits because they're so easy to understand—but they "don't do analogue." They barely understand Ohm's Law, and most of them don't even know what "impedance" means. So much for the great leap forward.

I completely agree that the best attitude in dealing with control-system design is that of the realist. You can't play the piano, either, if you're wondering if the keyboard is really there.

I don't think I agree with you about the gradual progression from closed-loop ("stiff" skating) to open-loop. But we can save a lot of detailed arguments if you will just get my book, *Behavior: The Control of Perception,* and go through it.

Concerning temperature control, I disagree with your statement that a thermostat must "control" room temperature if it is to control the measured temperature, the input. It must be able to increase and decrease the heat content in the room (that is what I mean by "affect" the room temperature), which in turn affects the temperature at the sensor. But the temperature at the sensor is affected by other things, too, not just by the furnace. If something cools air near the sensor, the average room temperature must be raised higher than the set-point in order to bring the sensor back to the set point. If you think in just a little more detail about physical processes in the room, I think you will see that there are temperature differences in various parts of a room, and that there are many variable sources of heat and losses of heat which alter temperature in various places in the room. Only the temperature of the sensor itself is controlled—that is, kept near the set point.

The realization that only the input is controlled by a control system (and not just the idea of feedback) was my "great idea." That's the meaning behind the title of my first book. You already understand control theory—just think about it, you'll get the idea. It's so simple that it took me three years or so to understand it. If you understand it in less time, that will show that you are smarter than I was.

Rick Marken: Izhak, if you know of any example of control by an open-loop system, then I want to know about it! And I would really like to know how it works. I claim that if a variable is demonstrably under control, then that variable is part of a closed-loop negative feedback loop. I think your examples of open-loop control are just higher-order variables under closed-loop control. For example, control is better in pursuit than in compensatory tracking. This is usually explained in terms of open-loop control—the person learns to "predict" the position of the target in pursuit tracking and makes "open loop" movements in anticipation. To the extent that the prediction is correct, then control appears to be better than in the compensatory case, where the temporal course of the target is invisible. I claim that the "prediction" is just a higher-order controlled variable—like a transition or sequence—and, to the extent that prediction works, it is really a result of the subject's ability to detect regularity in that higher-order variable. The subject is then controlling a transition, not just the instantaneous configuration of target-cursor discrepancy. The fact that this is what is occurring can be tested. If the prediction is open-loop, then failure of the prediction should result in no adjustment—after all, open-loop means not monitoring the consequence of the output. There are experiments using regular, predictable patterns of transitions as targets which show clearly that, when the pattern is changed (from a circle to another pattern), there is a clear adjustment to this change. The fact that the subject is controlling the sequence of transition rather than just the target/cursor configuration is evidenced by the rather long time it takes to makes the switch—about 400 milliseconds.

Still, I would be very interested in an example of what you consider to be open-loop control.

Izhak Bar-Kana: Rick, you are so enthusiastic about what you believe in, and work so hard on it! This is the way to results, even if at this stage or another one is wrong, and you are not wrong. Did you miss the start of my letter? I am a feedback-control guy, and I do not agree to give up the closed loop, not even for the sake of this discussion. I only claim that the dosed loop does not explain everything, and, I would dare to say, in particular when brain, learning, and intelligence are involved.

A curiosity. Recently I had an argument with a distinguished colleague, whose argument was: I know the desired output of my plant, and I know the plant perfectly (the transfer function); therefore, I can use this desired output as an input to something which performs exactly the inverse of the plant. (By the way, for everybody, the "plant" or the "process" is what we want to control, but is already there, given. For example, the motor, the ship, the tank. We can only add sensors and the "controller.") This way, the output of the plant will do exactly what I want. And then, I say: My friend, this is too idealistic, one can never know exactly the plant, the disturbances, etc. Furthermore, a transfer function followed by its inverse is not the equivalent of 1. (There is a problem of noncontrollability—I don't want to elaborate.) One cannot control without the closed-loop.

As you see, it is peculiar that I am now in the position to convince you that one can control without feedback. But now, let us go on from his idea: Assume, for a moment, that we know *exactly* the plant. We can then design a controller such that the plant will perform exactly what we want (not necessarily "exactly," but in a satisfactory way). If we do not know, let us try to learn about it, or to "identify." Once we identify, we can design our controller. Now, I ask, how can one be sure that the identification works correctly? And how can I be sure that the plant will not be destroyed before the identification is finished? And what if the identification process is corrupted by noise? All these thoughts occurred when I wanted to design some controller for a manipulator. It is a nonlinear system, and because of that, the load varies a great deal. If one uses only the high gains needed in difficult situations, one only amplifies noise in other situations. If one uses low gains, then performance cannot be obtained. Then, based on very little prior knowledge about the manipulator, I build a simple adaptive controller *(closed-loop)* such that the gains move up and down as a function of the tracking error, and the performance is quite good. However, since I do not use knowledge, the adaptive gains "work" very hard. Therefore, in parallel with the controlled plant, I use an identifier. The SAC (simple adaptive control) guarantees that no disaster will happen, even if the identification does not work properly. However, when the identification is correct, the controller based on it takes over and my adaptive gains (of SAC) decrease, and may vanish if they are not needed. The closed loop is there, and if the tracking error tries to increase, it will push it back. *Yet,* there is a signal, directly from the input reference to the plant, and for most situations, is now the only signal which controls the position of the manipulator. Is it not enough? Is it not capable of accounting for disturbances? Is it a bad control system? A bad control system is just that, it does not become a non-control system. And now, because my closed-loop gains are adaptive, even when they are called upon to correct for uncertainties, they do it at much lower values, usually, than before the addition of the open-loop control.

I wish I knew how the brain does its modeling and learning. I think that the identifier must use low gains for identification ("slow identification" or 'long-term memory") so that it is not much affected by nonrelevant transients, and only stores relevant knowledge. The SAC must be fast, to get the gain needed when it is needed.

Bill, now the argument is very close. What you call the input we should both call the "controlled variable." If I measure the temperature in the neighborhood of the sensor, then this is the controlled variable. However, this is not the end of the story. A good control system would be careful with its sensors and (as organisms do) would use some redundancy, measuring the temperature in various points, thus maintaining some relevance in the measurements, by averaging the various measurements or even by eliminating unusual inputs. Furthermore, a lot of noise might have been added to the measurement, and a close look would reveal a lot of filtering used to "ignore" this part of the input, and pass only the signal which is relevant to the controlled variable. Only a very primitive control system would just respond to any signal (or sensation) received as input, and respond to any spike of noise which comes from who knows where. There is indeed the danger that the control system might respond to any input signal, even if it is not related to the designed controlled variable, but a good control system takes care of it, sometimes using prior knowledge, sometimes identifying the disturbance and compensating for it, sometimes filtering undesired signals.

Maybe our different backgrounds lead us to put emphasis on various aspects of the same phenomenon. You claim that I control the input, because my output was disturbed by some bias (mainly constant disturbance), and I claim that I do not control the input just like that, because in most cases it is mixed with noise and must be processed before I can be confident (always only to some extent) that it represents the designed controlled variable. About open-loop control, the manipulator example above is also relevant. About open-loop in organisms, I don't have another example than myself, and that is only one sample. I know that before I know my car or my trajectory, I use a very stiff closed-loop system. If I get off the highway, my control loop takes me back immediately, and I might reach the other side or crash into the cars moving in the opposite direction. Experience, or the teacher, teaches me to ignore this signal, keep the same direction, and come back slowly. I see here a combination of closed-loop and open-loop control, but this is only one opinion.

Rick, you are right that the evidence on prediction does not prove anything about the existence of open-loop control. Actually, from an input-output point of view, open-loop and closed-loop schemes are equivalent. They have different properties with regard to sensitivity to uncertainties, stability (including oscillatory or non-oscillatory responses), etc. But the argument is not which concept is better, only whether open-loop control is there. In my case, of the manipulator, I know that the open-loop controller is there, because I put it there, and I see that the closed-loop "disappears" when open-loop is sufficient, because the gains vanish. I have no intention to claim that an open-loop system can deal with drift, disturbances, uncertainties, etc. I only claim that forms of open-loop control might exist along with closed-loop. The relative gains, or weightings, might vary, and in some situations, each one can be zero. If you have evidence that in biological systems there is no open-loop control, then I cannot argue, because I simply don't know. But one can explain some behavior either way, again, because from the reference-controlled-variable point of view, they are equivalent.

Otherwise, I can add to your argument: a good closed-loop control system does use prediction *in the closed loop*, if it is needed. Any phase lead in the forward (I am afraid to call it feedforward) path, or velocity (tachometer) feedback *along* with position feedback supply information on the future development, in other words, prediction. And this is only an elementary example of prediction.

Bill Powers: Izhak, I believe that one of the important steps one must take in understanding human control systems is to recognize the role of perception. If we were building an artificial device, then as its designers we would be aware of the true state of the controlled variable, and the state of the signal produced by the system's sensors as it represents that variable. We would be able to see what is signal and what is noise; we would know whether the signal is properly calibrated and linearly related to the external variable. But in the human organism, there is no third party who knows all this. I think that as a designer of systems, you can appreciate how the world must look to the system itself. It exists only in the form of the sensor signals, and whatever other signals are derived from sensor signals by computation. How can we discover the world as it is represented by the outputs of human sensors and sensory computers? The answer is so simple that nobody seems to have thought of it: just look around. Feel, taste, hear, and see. There it is. You are already experiencing the sensors' output signals, and you do not experience the world which gives rise to them. I do not think that this insight would come easily to anyone who has not worked with artificial sensors. But most people, in the end, understand it if they persist.

If you look at your experiences as signals in this way, you will realize that there is very little noise in them—they are almost perfectly noise-free. Only in unusual circumstances—near-perfect silence, the threshold of darkness—do we experience our perceptions as behaving in a way which seems at all "noisy." Also, you will realize that linearity and calibration mean next to nothing, because you are looking at the output, not the input, of the perceptual functions. You have nothing to calibrate them against but the outputs of other perceptual functions. You even pick reference signals from previously experienced sensory signals, so the scale of reference settings contains the same nonlinearities. You can say that *this* perception is nonlinear with respect to *that* perception or with respect to a meter reading (another perception), but you cannot say whether all perceptions are nonlinear in a different way with respect to their causes.

I think that this orientation makes a great difference in the way we build models of human behavior. We must realize that however the brain manages to bring external variables under control, it must manage this completely on the basis of information available to it through its senses—its uncalibrated senses. It cannot look at the plant (the universe outside) and see what compensations are necessary in order to represent its variables properly. It knows only the variables, and even then only after they have already been represented as internal signals. The only way it can identify properties of the plant is through experience with sensory representations of the plant's behavioral variables in relation to sensory representations of the organism's own output efforts. It knows something of the inputs to the plant, and something of the outputs from it, but it knows nothing directly about the plant. The world outside is a black box. We who have seen artificial control systems both from the outside and from the inside have some advantage in understanding this situation, because we can appreciate what is lost when you lose that disembodied vantage point from which you can see what is happening on both sides of the sensory barrier and on both sides of the output boundary.

This means that when we try to guess how organisms learn to do things like adjusting their internal part of the loop gain appropriately, as in your simple adaptive control, we must try to see how they can do it on the basis of information available *inside* the controlling system. In your case, you have accomplished that: just use the information in the error signal, which is inside the system. If the average error signal gets too large (which takes in many possible problems, such as oscillations), reduce the gain of the output function, the effector part of the system. I don't know how your identifier works—does it, too, work only on the basis of sensory signals available to the system as a whole? Or does it need external intelligence to tell it what to identify?

When we build models of human control systems, we naturally have to play the part of the "external intelligence" just to set up a plausible system. But our goal must be to learn how the system itself can come to acquire those design features we find necessary, without knowing what we know about our own created system designs. I have felt for a long time that the people trying to reproduce human "pattern recognition" have been on the wrong track, because a "teacher" is an essential part of their approach. Some external intelligence must tell the recognizer if it is right or wrong. Real organisms do not have such a teacher, not when it comes to learning the basic perceptual and control processes themselves. Recent work on neural networks and perceptrons is, I think, a little closer to the right approach, because the system in part creates its own organization. But there is still a teacher who knows the right answer. Real organisms never know if their answers are right, except in terms of how well they serve to control what happens to the organisms.

My reason for total rejection of open-loop control is based on thoughts like these. How can the organism find the feed-forward output signal which will create “almost the right behavior” of the plant, without monitoring the behavior of the plant? You have to imagine setting the direct output signal to just the right value which will keep the car centered on the road without *ever* seeing the relationship of the car to the road. It is impossible without the aid of some third party who knows what the plant is really doing. You could, of course, have sampled control, so that corrections are applied only now and then. But that is still control, and it cannot work without feedback. All of what is commonly called “feedforward” is really explainable only in terms of a hierarchy, of feedback control systems (in many cases, including those in an engineer silently standing by in the background, screwdriver in hand). Only through feedback can the so-called feedforward be properly adjusted. There is no friendly omniscient engineer in the background adjusting our own “feedforwards” for us.

Your example of overcontrolling a car is, I think, only a description of how we learn the right dynamics of response and get the control system stabilized. In the end, we have very fine high-gain control with proper temporal filtering so that oscillations are eliminated. But we do not notice this control because it takes place in our midbrains, brainstems, cerebellums, and spinal cords, where we seldom pay much conscious attention to what is going on. It doesn’t seem that we are exerting much control effort when we drive down a straight and level road, but just watch the steering wheel! It moves with every little bump in the road, every slight change in the crosswind, every little tilt of the roadbed. This control system is extremely sensitive to error—but it keeps the error very small, so it does not have to make big efforts. Unless, of course, there’s a big disturbance. We habitually observe from a higher level of abstraction, and we don’t notice the errors or the corrections because they are happening at lower levels. The car seems to go straight by itself. But just try holding the wheel absolutely still, and you will see that significant disturbances are always present. Their effects are cumulative. If their effects aren’t precisely corrected, the car will quickly go off the road.

Rick Marken: These are some comments on the thread related to open/closed-loop control. Bill has looked at the issue “from the inside.” Let me try it again “from the outside”—looking at control as the observer of a controlling system, rather than as an example of one of these systems. When we look at a living system, we see that it produces many consequences. These consequences are potentially variable—the temperature at the skin, the position of a limb, etc. The value of the variable at any time depends on many factors—the “causes” of the variable.

Thus, y = f(a, b, c, o), where y is the variable of interest and a, b, c, and o are the variables which "cause" y; f is the function which determines how y varies (over time) as a function of variation in the causal variables. If a, b, c, and o vary over time, then y should vary over time in a manner determined by f. If, however, y remains approximately constant over time, then we might imagine that something fishy is going on. Stability of y could happen by chance—the variations in a, b, c, and o could just happen to produce a constant y. But the longer this goes on, the less likely it becomes that stability is occurring by chance.

Moreover, if we can trace the stability of y to variations in o, which happens to be the causal influence on y exerted by the living system, then there is strong evidence that o is systematically counteracting the effect of a, b and c on y. I take these two pieces of evidence—the stability of y and the fact that system outputs are the sole cause of this stability—as evidence that y is a controlled variable. This is evidence of control "from outside the living system." It says nothing about how this control is achieved. What I am claiming is that the only organization we currently know of which can provide an explanation for control is control theory—that is, the theory that y (the controlled variable) is part of a closed negative-feedback control loop. One other part of that loop must be a reference signal which specifies the particular value at which y is stabilized. If y is stabilized at different values, then this reference signal must be variable. Observation of the environment of the control system reveals no variables "out there" which could possibly function as the reference signal (although people have been fooled into thinking that "targets" in tracking work this way; the simple way to show that they don't function as references is to show that people can reliably keep y stabilized at values different from the target value "when they want to"). Thus, the reference signal must be inside the system itself.

The controlled variable, y, need not be a simple aspect of the system's environment. We see living systems controlling very complex variables, such as their relationships with other organisms. I have seen people keep a variable called "in love" at nearly the same level for periods of years. The ability to control such variables implies an ability of these systems to perceive such variables. Perception, from a control-theory point of view, is not some arcane discipline of only peripheral interest to psychologists. Perception becomes central to understanding human nature. What people do depends on what they perceive and where they want those perceptions to be.

This model of what we see as the behavior of organisms is radically different than other models of behavior currently embraced by most life scientists. It is a model which works, which satisfies the requirements of scientific method, and which provides a comfortably humanistic view of human nature. That is the reason for my enthusiasm, Izhak. I guess I "believe" in control theory; but not in the usual sense of belief. I am not reverential toward it. If it proves to be wrong, I will happily abandon it for the improved point of view. I believe, based on experimental evidence and matching the behavior of models to that of living systems, that control theory currently gives us the best (and only) model of how people (and other living systems) work. I think this is not only scientifically important, but socially important. I think a case can be made for the notion that people have been screwing around with each other and making life more difficult for each other because they have been looking at each other as a particular kind of object—one which can be controlled from the outside. The control model shows that this is not only false, but also a sure recipe for conflict. And I think most people would agree that conflict between people (classes, religions, nationalities, whatever) has been the continuing obstacle to the possibility of every individual (other than the winners of each conflict) leading a graceful, dignified and satisfying existence.

Gary Cziko: An open question to Bill Powers (or other "serious" modelers): I just finished giving a presentation which involved showing Bill Powers' computer simulation of a simple control system. One person in the audience made the point that because the computer was doing the controlling, it had to be a iterative system. My somewhat lame reply was that, yes, it is iterative on the computer, but that the slowing factor added to the model makes it work like a continuous system.

But I suppose the point the person was making is that iterative control can work, in which case we do have responses which are computed based on the present static state of a number of variables. This is what the computer does, and I suppose all digital control systems as used in engineering do the same. *But* there is nothing in either the data we get from real subjects or in what we know about nervous system and muscle physiology which leads us to believe that control works this way in organisms. So we use the digital computer with slowing as an approximation to the continuous control we get with living control systems.

Somebody let me know if I'm on the right track here.

Bill Powers: Yes, you're on the right track, Gary. The slowing factor is introduced to keep variables (at least one variable in the loop) from jumping instantly from the old value to the next computed value. A real arm obviously can't be in one position at one moment, and in a position 20 degrees away in the next millisecond. The slowing factor is chosen to fit with the assumed physical time represented by one iteration so that the actual amount of movement is similar to the real amount of movement over the least element of time. The less the time represented by one iteration, the more slowly the variables must change. The slowing factor, being in the denominator, must increase as dt (the one-iteration time interval) decreases.

When we run models, we want to run them quickly, so we can try the model over and over while adjusting parameters for best fit with the real data. So we start with a relatively large value of dt. If the interval is too long, we don't get as good a fit as when it is shorter. At some length of time interval, around 1 /20 to 1 /30 second for some of the models I have made, making the interval shorter just slows down the computations without improving the model any more. This shows that over roughly 1 /30 second, the variables in the model vary slowly enough so that the response is essentially the same as if the sampling were infinitely fast. We choose the interval dt so the results are the same as if we were sampling the behavior at an infinite rate.

Even with this explanation, there is still often a problem in getting people to see the difference between an iterative quasi-analogue computation and a sequential computation. In a sequential computation, each variable is calculated in turn, just as in our computer simulations. But the mental image which the listener is thinking of is really cast in terms of *events.* First there is an input event which causes a perceptual signal event. Then the perceptual signal event is compared with the reference signal to yield an error signal event. Then the error signal event causes an output event—a response. And while these events have been taking place, what has been going on at the input? This is the question they overlook; they assume that the input event is finished, so nothing will happen until the next input event occurs, perhaps "triggered" by the response. So each function in the loop takes its turn in acting, and then lies quiescent until it's aroused again. It's never aroused again before all of the other functions have had their turns.

In the real system, of course, the input varies continuously. All of the functions are doing something all of the time. There might be a delay before the next function in the loop receives a given input value, but during the delay the input continues to change. So the next function receives a continuously changing signal, delayed, even while new changes are being introduced at the input. There is a pipeline effect. It's like talking to someone over a satellite link. Your voice vibrations are received at the other end continuously, but delayed by the length of the link. This is very different from thinking about input events and output events.

A truly sequential system would be represented by a feedback loop, digitally calculated, without any slowing factor. We can boil such a loop down to an extremely simple example:

A = B
B = -l0A

If you start with any value for B (except exactly 0), this loop will run away on successive calculations. But suppose we introduce a slowing factor:

A = A+(B-A)/S
B = 10A

Now the loop will converge so that both B and A approach zero, provided that S, the slowing factor, is larger than 5.5. If S is 11, the final state will be reached in one jump. If it is larger than 11, the approach to the final state will be monotonic from any starting value of B.

Even the cybernetician William Ashby fell into the trap of sequential calculation. He concluded that negative feedback systems couldn't have a loop gain of -1 or more negative and still be stable (note that the above system has a loop gain of -10).

The implicit reference signal in the equation above is 0. You can put in a nonzero reference level for A in the second equation:

B = 10(A* - A), where A* is the reference level

Now the system will approach a state with A nearly at the value A*, from any starting condition. You can make A come closer to A* by raising the loop gain:

B = 100(A* - A)

But the system will oscillate unless you increase the slowing factor. If S is made equal to 101, the final state will be reached in one jump. If S is larger than 101 (say, 300), the approach will be monotonic. If G is the loop gain, then S must be greater than or equal to G + 1 in order to get a stable approach to the final state. Note that G is a positive number for negative feedback because we are subtracting A in the above equation.

The final state you reach is predicted by solving the *first* two equations (without the slowing factor) as a simultaneous pair. If there is a non-zero reference value A*, solve this pair of equations:

A = B
B = G(A* - A)

It's not obvious, but introducing the slowing factor converts the pair of equations from a simple algebraic system into a differential equation.

That's why we are able to stabilize its behavior in time, even with loop gains as large as we please.

I realize that you're not going to take a naive audience through all of this in a one-shot lecture. But if you play with these equations enough to get the feel of what is going on, plugging numbers in and running the iterations, you'll probably be able to cope with the misunderstandings a little better.

Izhak Bar-Kana: I assume that Bill's equations with S are as follows (the minus was missing):

A = A+(B-A)/S
B = -l0A

Now, the program takes it as written, but if we want to describe what is going on, the equations should be these:

A(k) = A(k-1) + [B(k-1) - A(k-1)]/S B(k) = -10 A(k)

Substituting B(k-1) in the first equation gives the following:

A(k) = [(S-11)/S]A(k-1)

We get the condition S is greater than 5.5, because we want the loop gain K = (S-11)/S to satisfy the stability condition |K| is less than 1. So, what is Ashby's mistake? What do you call the loop gain?

In general, we get these equations:

A(k) = A(k-1) + [B(k-1)-A(k-1)]/S
B(k) = G[A* -A(k)]

Notice that A(k) and B(k) are Anew and Bnew, while A(k-1) and B(k-1) are Aold and Bold. A* is constant here, but it might be any function of time. The same substitution gives this result:

A(k) = A(k-1) - (G+1)A(k-1)/S + GA*/S

or

A(k) = [1-(G+1) /S]A(k-1) + GA*/S

Now, one selects G and S so the loop gain is |K| is less than 1, where K = 1 -(G+1)/S. If this condition is satisfied, and thus a stable equilibrium point exists, it is reached when A(k) = A(k-1), and we get the equation:

$$A(k) = [G/(G+1)]A^*$$

(if I don't have an error of algebra), which tells us that in such a simple system one cannot have perfect following even for a constant input (sorry, I mean reference), unless G is infinite.

Bill Powers: Izhak, your analysis is precisely the same as mine, and you found a shorter way to prove that $S = 2/(1 + G)$ is the minimum value of S for convergence (see my 1978 article in *Psychology Review* for a longer way).

Why can't I learn to get critical signs right when I publish equations? You are correct about the sign of "10B," of course.

Actually, with S greater than 5.5 but less than 11, the approach to the final state is oscillatory, and the oscillations are an artifact of calculation (if you're trying to model an underlying continuous system). The oscillations occur at the iteration frequency and are not tied to physical time. Only when S is greater than 11 can you model the real motions of a physical system.

You have defined the loop gain here a little differently, so that it is the gain allowed by the slowing factor on each iteration. I wish I had thought of that—it's so easy. I would call the loop gain G (or 10 or 100, depending on which equation you read) because that is the gain which predicts the limiting case (infinite integrations)—that is, A[infinity]. In the limit, $A = G/(1+G)A^*$, and S drops out. You arrive at the same result, quite correctly, by specifying that A ceases to change. The same result is given by taking the equations:

$$A = B$$
$$B = G(A^* - A)$$

and solving them simultaneously by substituting B for A in the second equation:

$$B = G(A^* - B), \text{ or}$$
$$B(1 + B) = GA^*, \text{ or}$$
$$B = [G/(1+G)]A^*$$

Solving these equations simultaneously is the same as saying that these two relationships hold *at the same time,* so this is a control system with zero time-lag and zero slowing. I use this as a way of showing that a control system which is properly stabilized behaves (in the steady state) just like a system with no lags. Of course, its dynamics will be different, but when you're interested in an overall view of relationships among variables in a control system, dynamics aren't the main subject.

As to where Ashby went wrong, he didn't use any slowing factor in his equations. Of course, when he set the loop gain to any number greater than -1, the system simply went into ever-increasing oscillations. From that, he concluded that negative feedback can't work with loop gains more negative than -1, and therefore that negative-feedback control must be very weak. Maybe that's why he gave up on the negative feedback model and used an open-loop compensation model instead. Ashby was a psychiatrist, after all. He didn't really know much about control theory.

You note that "in such a simple system one cannot have perfect following even for a constant input (sorry, I mean reference), unless G is infinite." Technically, you're correct. But practically, with a G of 100 or 200, the system will keep errors small enough to be ignored in models of behavior. The actual values measured experimentally for subjects in tracking experiments come out in the range from about 50 to 200. So if the model's G is set too high, it will behave too perfectly. With the correct G, the model will make errors similar to those which the subject makes. We have taken to using an integration factor because, with gains that high, there is no significant difference between a pure integrator and a high-gain proportional system with an appropriate slowing factor. I went through a comparative analysis a few months ago and satisfied myself of that. When you're retired, who else do you have to satisfy? And it's all right if you say "input" here, because in the context we will all recognize that it means "reference input" and not "sensory input."

I *hope* you didn't make any algebraic errors, because the derivations looked fine to me. I don't usually bother with the subscripts, but your use of them is the same as mine when I put them in. When one does most calculations through programming, an equal sign comes to be understood as the replacement operation. Bad habit, no doubt.

Izhak Bar-Kana: I cannot become one of the family on the Net, especially when the discussion becomes philosophical. I can smile when you give me the thermostat example as a living illustration for the control of the input, because under the conditions you describe, I would fire the designer. I am not sure I know where I belong, because I try to get something from everybody, so I try, at least, to read your discussions. One thing I do know: I am an engineer—I would say a bloody old engineer—and cannot change overnight.

For me, there is no reason for existence of any control loop or, better, control system, if it does not control the output. I must use measurements to monitor this output, and if I am wrong, I might end up controlling something else. But, in the same way, the control signal I design is going to be transmitted through sore actuator, and if I am not careful, it might become very different from the desired control signal. If there is a danger that the input I measure does not represent the output in an acceptable way, I will use lots of filtering (estimation) or lots of redundancy. It might be that in organisms, the emphasis is on the other aspect; I don't know. However, if a simple engineering system includes so many redundant loops, I have the feeling that the extraordinary redundancy in living systems has the same role: to avoid controlling the measured feedback input or responding to a measured reference signal which does not reproduce, in a reasonably exact way, the realworld external values.

But I see I am getting philosophical without even trying. All I wanted to say is that G is not the loop gain, once you use the "slowing factor" S. The loop gain is now given by K = 1- (G+1)/S, from this equation:

$$A(k) = [1-(G+1)/S]A(k-1) + GA^*/S$$

and it *must* be less than 1 to get a stable system. Of course, if K is less than 0, A will change signs every interval, and in a first-order system (with only one delay involved) this oscillation can be prevented using K positive. High gain is a solution when noise is not involved, otherwise the difference between integration and high gain becomes evident: while the high gain amplifies any noise, integration averages it.

Gary, yes, the digital computer is a very easy and handy way to approximate and simulate continuous systems. When the continuous system is sufficiently slow and the sampling is sufficiently fast, one can ignore the difference. In more complex cases, there is an entire theory about how to switch from the continuous to the discrete domain and vice versa. This is not a trivial problem. There are phenomena which cannot be exactly reproduced in the discrete simulation (what happens at collisions, etc.). When one wants only to simulate an approximate behavior, especially in closed-loop, many parameters can change without affecting the results, and any discrete approximation will do.

Advantages of discretization? It is so convenient! Try to implement a slow process with a time constant of, say, 10 minutes. One might need the earth for the capacitor which would be required. In discrete form, it is just a line of code. But most important is implementation of timevarying and nonlinear parameters and algorithms which are almost impossible in analog form. And, by the way, delays. Very fast processes, however, cannot, or cannot yet, be implemented digitally, and analog circuitry has made some progress. So, actually I see a combination of both as the future solution for computation. Complex simulators use "hardware-in-the-loop": those parts which are too fast or cannot be simulated with confidence are used directly in the loop, of course using D/A (digital-to-analog) and A/D converters. This brings us to real-time simulation, which is another opera.

I used to simulate very complex systems, such as planes, flexible structures, etc., with large ranges of time constants. There are simulation languages which allow you to write the equations of the continuous systems. The translation to the discrete world is done by the computer, sometimes using different time intervals for different integrals, so the errors are maintained below some admissible value. In these cases, the precision is almost continuous.

The slowing factor does not make it work like a continuous system, it only makes it work. This is also the danger of simulation, especially when presented to inexperienced students. They take the results for granted, because “the computer shows.” But the computer shows exactly what we supply it with. As I understand it, you do not have any detailed models of the various components which together form the simulated closed-loop. In this case, one must emphasize the fact that by using a simple model, one manages to reproduce the behavior of the real thing, to some extent. But not vice versa: the real organism does not behave this way, just because the computer shows.

Bill Powers: Izhak, I talk about the output of the control system, which is the input to the plant (the environment), while you talk about the output of the plant, which is the input to the control system. We don’t control the input to the plant—that is varied as disturbances require, so the state of the control system’s output is just as unpredictable as the disturbances are. The output of the plant is under control, and so is predictable. That is the same as saying that the input to the control system is under control: the only difference between saying “input” and “output” in that case is whether you take the plant’s point of view or the sensor’s. You see that I am separating the control system from the plant that is controlled; perhaps you draw the boundaries differently.

In artificial control systems, the engineer can see both the sensor signal and the objective variable to which it corresponds—what you call the output (of the plant). In living control systems, the observer (the one which matters) is inside the system and can see only the sensory input. The variable in the plant (the environment) can only be inferred; it is not available to direct inspection by the control system. This makes a great deal of difference when you are talking about systems which, in effect, design themselves.

In speaking of artificial systems, it is optional whether you consider the controlled variable to be an input or an output variable: it is the same variable in any case, just outside the sensor. In speaking of living control systems, however, where we must account not only for their operation but for the internal organizing processes which bring them into being, we must choose the "input" interpretation. In fact, we must say that the perceptual signal itself is really the controlled variable, for sensors can vary in their properties.

When the sensor's calibration changes, the perceptual signal remains under control in the same state as before, but the external variable on the other side of the sensor is brought to a new value by the control system. If we understand that the perceptual signal is the controlled variable, then we can understand how the behavior of the system changes when its perceptual systems reorganize. If we focus on the external processes alone, we will see only that something has disturbed the control process, thrown it out of kilter. We might even conclude that it has failed, when all it has done is to change its definition of its environment, possibly by mistake, but also possibly for its own purposes.

So I think that we have to think of control as control of input, if we are to grasp what is really meant by saying that we, ourselves, are control systems.

Izhak Bar-Kana: Bill, you seem to repeat some arguments I was trying to use when I started participating in this Net, and I felt that we used the same names for different things, and vice versa. But then I came to understand from you, and even more so from Rick Marken, that things are much more profound.

I am sorry, but for all of my (engineering) life I have been used to "input controls" and "output is controlled." The "control system" includes everything, and, of course, the plant. The part of the control system which controls the plant is the "controller." Now, if the input to the control system (to the controller, and through it, to the plant) is not zero, it will affect the plant. If it is zero, it will not affect the plant. In a closed-loop system, the input is obtained by the difference between the reference input (in tracking systems, control systems with feedback gain one, it is also the desired output) and the measured output. It is clear that the control signal, the signal which affects the controller and the plant, can be only the measured value of the real signal. Similarly, biological systems can only use the sensorial perception as the *control signal,* to affect their control system, and all of the various stages and values, up to the value which is called the "controlled variable." I really don't understand why this language, which you do know, had to be changed in such a way that Rick Marken cannot even talk to someone I call a control guy, because the old-fashioned engineer cannot accept the idea that any control system controls its input.

Now we are in a closed-loop, and you can again change the order.

I regret that it separates you from the general family of control research. More so, since we want to learn about the behavior of organisms from psychologists. I can ignore the linguistic differences, or at least try to, and try to get the ideas, because I don't know a better group and discussion. But why do you have to speak French in the middle of English?

The most intelligent system I might dream of designing does not come even close to the simplest organism. In my humble opinion, again, one of the reasons for the huge and not always motivated (apparently) redundancy in organisms is to prevent an occasionally wrong measurement (or input feedback) from replacing the correct output the control system is meant to control.

Tom Bourbon: Izhak, you have the respect of those of us who labor to understand *living* control systems. I am certain your life would be simpler were you to decide that we are a bunch of misguided nuts who cannot get our control-system diagrams and labels right!

Perhaps I am wrong, but part of the problem when you speak of engineering (designing and building) a mechanical control system and we speak of trying to describe and explain the control created by the living things we find already acting in the world is that we can't design and know all about the living systems. What is more, the variety of control theory we are trying to develop must compete with a host of already established and widely believed theories and disciplines, so we must direct most of our effort to persuading followers of those disciplines that there is even anything out there to notice which is different from what they already know. It is unfortunate that, in the process, some of what we say seems wrong to the part of the engineering community which is probably closest to us.

Living control systems were not designed by us: we found them inhabiting a world which had already buried them in a host of sciences and disciplines—the life sciences, social sciences, and behavioral sciences, recently joined by the neurosciences, cognitive sciences, and many, many more. For the most part, the practitioners of those disciplines and sciences do not recognize that living systems control *anything.* Rather, they speak of the behavior (actions) of living things as *controlled by* antecedents, whether from the environment (e.g., stimuli, contexts, gods, societies) or from somewhere inside (e.g., mind, soul, schema, plan, commands from the motor cortex). They invoke linear causes, and they reject control by living things.

All a living system knows of "the world" are its own sensory experiences of the world, so it follows that all a living control system can control, from its own perspective, are its own sensory experiences. And there is abundant and conclusive evidence that sensory experiences do not correspond *directly* with the environment. Perceptions as simple as those of brightness, hue, loudness, heaviness, and the like reflect states of the perceptual apparatus of the organism—*directly*—and they *fail to correspond directly* with any unique state of the environment. The state of adaptation of sensory receptors, the surrounding stimulus field, the relative sizes or magnitudes of different elements of the stimulus field, the relative temporal durations of stimulus elements, and many other variables can combine in many different ways to produce the *same* perceptual experience. Hence, a person, like any other organism, can have the same perceptual experience in the presence of a nearly infinite array of different combinations of elements in the environment and in the organism's own physiology.

Because perception does not correspond one-to-one with any unique state of the environment, it follows that an organism which acts to control its own perceptions is not controlling a unique state of the environment, hence is not producing (controlling) a unique state of its actions (output). The specific actions of the organism and the remote environmental consequences of those actions can vary dramatically, yet the organism experiences uniform perceptions. And it is certainly true that an organism which always produces the same actions and remote consequences in the environment will experience variable, not constant and controlled, perceptions.

In your engineering applications, zero input (by that do you mean zero perceived error—a state *internal* to the organism?) leads to zero output. But an organism which adopts a new reference to experience an absent perception experiences zero perceptual input, which creates in the organism a non-zero error, which drives the behavioral actions (output) of the organism to create the desired perceptual experience, which does not, for a perceiving organism, correspond directly with an objective state of the environment. A bird with a reference to sense a not-yet constructed nest experiences zero perceptual input of nest, and it acts until it experiences that perception. And a sculptor who decides to sculpt a bird on a nest experiences zero perceptual input and acts until that experience exists—whether any other person recognizes the finished sculpture as bird-on-nest, or not. To the artist, that is not important (not even if the artist must sell the sculpture to buy food—all that matters is that someone else desires the sculpture and pays what the artist asks).

Those are the kinds of control we find in the world of living control systems. The best we can do is look for situations in which the variables through which the organism or person achieves its control of perception are also sufficiently stable from *our* perspective that we have a due as to what the organism or person is controlling—in its own perceptions. Certainly, the one we observe is not controlling our experience—not as its primary goal.

Interestingly, it is true that the category of humans known as controlsystem engineers do enjoy a privileged position relative to the control systems they design, construct, and study. They do know the references and the "objective" states of the relevant variables in the environments of those systems. In fact, what the control-system engineer intends is that his or her perceptions of the states of those variables in the environment of the artificial control system will match his or her chosen reference. In that context, it is easy to understand why the engineer would speak of the artificial system controlling *its* output—what the artificial system represents is a way for the designer and builder to control *his* or *her* perceptions, relative to her or his references.

I do not know if any of this helps, Izhak. If anything I say violates too many of your ideas about control processes, please tell me.

Rick Marken: Izhak, the difference between input control and control of input is not just a language difference—it's the whole point.

For an engineering psychologist, the organism experiences error due to the discrepancy between an objective reference and input event (the target and cursor in tracking; sometimes the error itself is considered the stimulus). For a Control Systems Group control theorist, both the error and the reference are *inside* the organism. The reference can be adjusted by the organism (by higher-level control systems) so the organism determines what constitutes an error; the organism is in control of the environment, not vice versa—a rather significant difference. The difference accounts for the appearance that organisms can voluntarily change the value at which an environmental input variable is controlled—it's as though the thermostat suddenly decided to keep the room at 65 rather than 72. This is the phenomenon which control theory is trying to point to: voluntarity or, better, purposefulness.

The controlled environmental variable is probably what you call the output which is controlled by the system. That's fine—but, of course, it is this output *as perceived by the organism* which is controlled, not the output itself. With organisms, there is no independent means of checking the validity of the perceptual representation of the environmental variable being controlled—all the organism gets are perceptions of the environment. We cannot look past our perceptions to see if we are controlling what we intend to control (as you do when you design a control system and make sure that it is controlling what you intend for it to be controlling; you can look beyond the sensors, the control system itself cannot). So, for a living control system, reference states of perceptions (not environmental outputs) *are* the intended ends of control actions.

Note, by the way, that the mathematics of the engineering psychology and CSG psychology approaches to control are nearly the same (at least, control works in both cases). The difference is in where you put the variable r (the reference signal). That's all there is to it. Small step for control theory; giant leap for understanding the nature of living systems.

Wayne Hershberger: Izhak, the yoked terms *cause* and *effect (prod* and *product, independent variable* and *dependent variable,* and *controlling* and *controlled),* respectively, have gotten linked to *input* and *output* in engineering, just as they have gotten linked to *stimulus* and *response* in psychology. This linkage goes way back—long before the development of control theory. Relatively recently, Ben Franklin "mislabeled" the polarity of electrical potentials, and just as engineers continue to use Franklin's terminology as an acceptable convention, even though that convention has misleading connotations, so they can and do use the cause/input and effect/output conventions, even though that convention also has misleading connotations.

Consider the following bizarre statement: The output of a furnacethermostat system is an input, not an output, but this input is an output, not an input. Although this grammatical sentence is *not* nonsense, it is certainly gibberish. Deciphered, the sentence reads: The output of the thermostat-furnace system (i.e., what it produces or does or controls) is a particular value of temperature sensed by a thermocouple (receptor input), not a particular amount of heat emitted by the furnace (output), but this sensed temperature (or receptor input) is an output (i.e., the dependent variable controlled by the system), not an input (i.e., it does not cause the temperature being produced).

As you can see, some truths *cannot* be expressed in engineering's input-output terminology without sounding ridiculous. Conversely, some things which are truly ridiculous can sound very true. As long as input means in and output means *out,* it is unseemly to use those same terms (input and output) also to denote cause and effect (or controlling and controlled). As you can well imagine, when one pairs the term input with cause (or controlling) and the term output with effect (or controlled), many of the connotations are as phony as a three dollar bill. I want no truck with that currency.

I suspect that you are not bothered all that much by these mischievous verbal connotations because you think mathematically most of the time (Franklin's convention poses no problem for me when I use Ohm's law, but it is a veritable pain in the ass when I try to think about PNP versus NPN transistors). Perhaps you can understand why those of us who wish to describe control systems in *English* can ill afford to be encumbered by blind (and blinding) anachronistic language habits developed in the context of a technology devoted to *calibrated* "control" systems (wherein the input/cause and output/effect convention is not problematic).

I am not for a moment saying that control engineers do not know what they are talking about when they say that control systems control their output. What I am saying is that engineers do not realize the mischief they are making for the rest of us when they use the terms input and output in this way. I am speaking for myself, but I think the other CSGers would agree.

Izhak Bar-Kana: I think most of you ignore some of my words.

One problem, for example, is that we here must control the position of a satellite up there. So, the position of the satellite is *the* object of the control. I have no direct means to measure it, as I have no direct means to affect it. Yet, everything in my control problem revolves around this control object. I can only use some remote measuring instrument, but I have no intention to control its output. Actually, one must learn how to ignore some of the input signals, because they are disturbances or noise. No one can convince an engineer to accept the idea that he and his control system control an input. You all seem to accept the idea that there is no control unless there is intention to control. Therefore, the position of the satellite will be monitored by multiple sensors, well-filtered and processed (to estimate some of plant's states, such as velocities, etc.) in such a way that the controlled variable remains at the value I am interested in, namely the position, or more generally, the motion of the satellite.

If the result of the multiple measuring, filtration, and processing is what you would call perception, I have no argument.

Tom, many thanks, but please do not *respect* or *trust* me. The only principles I trust are: 1) never believe in principles, 2) no one here is God, and 3) no one is really dumb. In spite of the fact that you, Rick, and Wayne seem to agree about the contents of your messages, and I think I have no problem with most of them, there is a difference in tone between your explanation and Rick's.

There is a big difference between Rick Marken when he clearly presents an idea and Rick Marken in an argument. Sometimes, it seems that the second has had *The Revelation,* or even touched *God*, and to hell with the others. Rick, I am afraid that when you talk about control, you have a steady-state image in the back of your mind. I *must* keep a dynamic image in my mind, because some of my sad experiences show that the steady state might be beautiful, but it is never reached. I must use some mathematics because many great, ingenious, and intuitive ideas proved wrong. It is easy to show that "if this is so, the gain must be so and so, and if the error is so, that lets us adapt the gain to be so and so." These arguments convince, and engineers like them. However, after a plane crashes, and one analyses (very difficult math, particularly in nonlinear systems) the aftermath, one discovers that things became unstable just because one used non-constant gains, even very carefully, namely within the "admissible" bounds which were tested with constant gains.

When I control the motion of my hand, this is the intention of the control, the object of the control, and I think that the corresponding control system controls the position of the motion of the hand. If you agree that there is intention in control, then this is the only intention. The fact that I must *measure, sense,* or *observe* this motion is a problem, not a principle. As you have observed lately, a closed-loop system is so built (if correctly built) so that the gain of reference-to-output is almost *one,* while the gain of disturbance-to-output is almost zero. We can show that an integration in the forward path makes the corresponding ratios actually 1 and 0, at least for constant reference inputs.

I try to stay aside when psychology is talked here. If you consider that the reference inside the organism is a great idea, showing that life affects the environment, and not vice versa, I am excited by that idea. Unfortunately, I cannot claim the same thing about the artificial control loops. I mean the environment makes a lot of trouble. I only have my own organism to observe, and personally I think that I can decide to drive here or there only after a long period of learning, and that the reference points within are a good mapping of the reference points out there. I might decide just to follow an internal reference with no relation to the outside word, but usually I stop after the right number of glasses. Furthermore, the reference command to be followed I would rather call *decision* than *control.*

If I want to move my hand, or a robot arm, I apply a force. If I meet resistance, I use more force. If there is an egg there, I must behave in a different way, and this is first of all a decision problem, or a detection problem. I don't call everything control theory. Many control people (engineers or not) do not know detection or decision theory, and they have to rediscover it again and again when it is needed—not the best way.

Wayne: Time flies like an arrow, fruit flies like a banana. I admit that this is not related to our topic, and I admit that what you call gibberish is gibberish indeed, but I cannot see how you relate it to any engineering. Still, I am happy you do not blame the Original Sin on me. I will never tell anyone he is wrong because his arguments remind me .... I can argue about right or wrong, understanding or not understanding, and so on. Too many arguments here blame me for talking like the behaviorists, like Wiener, calibrationists, and who knows what. I think you have better arguments for your position than calling the control of plants a simple misuse of words, even if control of perception is needed and correct under your paradigm.

What an engineer means is that his system controls the position of the plane, and he calls it an output. It has nothing to do with any old-fashioned calibration, as it has nothing to do with the Middle Ages. He will do anything that is needed, and possible, to make this plane follow the desired path. It has nothing to do with grammar, nor with the fact than any input is an output of something else and vice versa. But I think we rotate now in a closed-loop with no reference whatsoever.

I first thought it was worth understanding your terminology and to bring it to some common denominator with the large family of control theory, but it is not very important. When I say "control," I have a plane or a robot in my mind, not a differential equation. My only problem is that this robot should be at a given position at a given time, no matter how I monitor (sense) its motion.

When you see these lines on the display, the desired output of so many control loops, all designed to satisfy your finest perceptions, even if they cannot control (or because they cannot control) your perceptions, if you can claim that an engineer does not care about what is input or output, I can only ask: Who do you call engineer?

Bill Powers: Izhak, it isn't that we ignore your words: it's that we can accept them as truth, but truth of a kind leaving out other important truths, particularly the one we have found the most startling and the most informative: the truth that a control system can control only what it senses.

In the world of engineering, the engineer has full knowledge of both the environment of a control system and the internal design of the control system. So he can point to a consequence of the system's actions and say, "There, that's the output which I want to be controlled, and here, in the system, is the feedback signal which represents that output." In doing this, he does not have to pay any attention to the fact that he must use his own senses to see that output. Literally, however, *for the engineer,* the output being controlled is known only in the form of a perception (whether aided by instruments or not). That is a fact, but it is irrelevant in engineering.

It is not irrelevant in trying to understand how the engineer works. When we look at the design of the engineer himself, according to our best neurological and physiological models, we can see that the engineer's entire world must exist in the form of sensory signals and higher-order functions of them, also represented as signals. In a way, you have given a nice example of this in talking about controlling a satellite: "So, the position of the satellite is *the* object of the control. I have no direct means to measure it, as I have no direct means to affect it. Yet, everything in my control problem revolves around this control object.

I can only use some remote measuring instrument, but I have no intention to control its output." So how does the engineer know of the position of the satellite (other than by looking up)? Only, as you say, by using a remote measuring instrument. He has some moderate amount of faith in the instrument, after calibrating it, but that does not change the fact that *all* he knows of the satellite's position is in the form of this instrument's reading (which consists of numerical digits, not positions). He does not, in fact, know the position itself. He knows only the reading, and he has a complex theoretical structure in his head which converts this reading into a concept called "position." He calls this concept the "output" he wants to control, but in strict literal truth, it is a perception.

The engineer might have no intention of controlling the output of the measuring instrument, but in fact that is all he can control. He has no other way of knowing the position of the satellite except through the use of earthbound measuring instruments (his eyes among them). He trusts that the instrument readings correspond in a regular fashion to the "actual position" of the satellite (with all necessary corrections applied, for example the time-lag of light rays and radar pulses, and the various motions of the earth itself). This trust is an epistemological statement, but its truth or falsity do not matter here: we are talking about practical requirements. The engineer *imagines* that he is controlling the position of a real satellite, up there in the sky, and he can produce all sorts of justifications for accepting this imagined correspondence. But he can't know that position without using the instruments, and he can't know the effect of his remote-control actions until he sees what the instrument readings do. Whatever he believes is actually going on, he is stuck, as a practical matter, with controlling a perceptual representation, not the thing itself. His epistemological beliefs make no practical difference at all.

All animals, and most human beings prior to the age of higher learning, necessarily act from the epistemological position that the perceived world is the world itself. There is, of course, no alternative. Speaking for human beings, the reality we know as solid and real, upon which we act and which we intentionally alter in some regards, is the only world there is. "Perceptions" don't exist except, for some of us, as philosophical abstractions or "signals" in a model. When we forget about models and philosophy and just look around, we see the world, not perceptions. When we look up into the night sky and see that serene and untwinkling point of light moving steadily and silently among the stars, we say, "Look at that! There's a satellite!" We don't say, "This is a perception of a satellite, a signal in my brain." When we point at the satellite, we see our own hands with forefinger extended. We don't say that there is a perception of a forefinger, nor do we pause to wonder about the relationship we call "extended." The relationship is just as much "out there" as the finger is. It would never occur to us to wonder what sort of thing it is, out there, which we call a "relationship."

This simple and self-evident world has conceptual holes in it. The biggest hole lies between the intention of pointing at the satellite and the immediately experienced actuality of pointing at the satellite. How is it that a mere intention, a figment of the mind, actually causes this pointing to occur? Control theory provides a plausible way to fill in this gap, a way which is as self-consistent and as consistent with observation as any finding of physics. In a manner of speaking, it is a finding of physics. To construct this model, however, we are forced to readjust our conception of the whole apparent reality, because the control model can work only if the satellite and the finger exist for the controlling system as signals produced by sensory inputs and subsequent computing functions in the brain.

In your objections to the concept of control of input, you have consistently assumed that the engineer can know the actual state of the output. Within the boundaries of the usual world of engineering, observing is not a brain process: it simply consists of noting what is there, while the role of the engineer's brain in making this possible remains silently in the background. In our explorations of control theory, however, we make this brain-in-the-background explicit. Even in talking about artificial control systems, we habitually take the point of view of the control system, something which few engineers would see any reason to do. We say, "If I were that control system, what world would I be experiencing?" And the answer, of course, is that the world would consist completely and exclusively of the signals coming out of the sensors. We could not know what is causing those signals; we could not even know whether they represented light or magnetism or sound. They are just signals. They get fancier labels only in the context of other signals which are also just signals—or in the mind of the engineer, who occupies an omniscient position in relation to this tiny control system and its surroundings.

When I speak of what "we" think on this subject, I am speaking of those who have internalized this model to the extent of relabeling their own ordinary experiences as "perceptions," at least when thinking in the modeling mode. This relabeling has come to most of "us" in a moment of sudden illumination which forever alters how we understand nature and ourselves. Nothing is changed in ordinary experience: "out there" still seems to be where it has always been. What changes is its meaning in relation to how we interpret the behavior of others and ourselves. This threshold of understanding is either passed or it is not. Once it is passed, the world of experience not only contains new implications, but *it makes a great deal more sense than it made before.* This is what has attracted so many people to the CSG version of control theory in the context of living systems. So many questions are answered, even those we hadn't thought of asking. So many holes are plugged which we hadn't even recognized as holes.

"Reality Therapy" and "Experiential Therapy"

Rick Marken: To Dag Forssell: I have your "Alignment/ Mission Statement" and "Discussion of Issues and Control Theory." The first seems to be a template for a statement of agreed-on higher-order goals for two control systems (people) working as partners in an engineering firm. It looks OK to me. I have a bit of trouble with phrases like "accept responsibility for our lives" and "efficient perception of reality." I also think the statement that control theory views people as controlling themselves misses the point by enough to be misleading.

Your "efficient perception of reality" statement makes me wonder—what' did the guy who developed "reality therapy" see in a model of behavior as the control of *perception?* Is the idea of reality therapy that *reality is perception?* If so, why use the term reality? It suggests a therapy which helps you get in touch with reality, which suggests that the therapist knows what reality is, and you (the therapee) should too. If I understood behavior as the control of perception, and problems requiring therapy as the results of perceiving things in ways which prevent non-conflicting control of those results, I would never have thought of calling my therapy "reality therapy." Maybe "control therapy" or "perceptual reconciliation therapy;" or, best, "conflict resolution therapy." But "reality therapy"? What could be more misleading? Clinicians, could you tell me why William Glasser, who claims to have understood control theory before he even discovered it, called his approach to therapy "reality therapy"? I smell condescension here.

Your discussion of control theory seems reasonable. It does emphasize the control of perception. I would suggest that you make clear the relevance of perceptual control to the problem of conflict and how to resolve it. After all, I think that's what the value of control theory is for effective management: finding ways to perceive the production process so that there will be minimal or no conflict between the cooperating contributors to the process.

Dag Forssell: Rick, thanks for reading my papers. I am glad you did not find any major misstatement on my part. I cannot ask for any thorough critique at this stage, since my presentation is not finished. The particular papers were extracts and summaries, respectively. I am trying to introduce control theory to industry in a fruitful way.

When you came across the word "reality," your configuration references immediately associated with "therapy" and "Glasser." I am *very* glad that I found the Control Systems Group, even if it was by way of Glasser's writings. I am not a student of Glasser any more (I can still see value in many of his writings, both from a medical perspective and in the more recent musings on quality, but he is dangerous because he totally misrepresents—*grossly* oversimplifies, apparently because it reduces his error signals—control theory and what one can learn from it).

Anyhow, upon closer scrutiny, you will note that the word "therapy" is nowhere to be found. I made reference to Abraham Maslow's admittedly unscientific observation that the most outstanding common denominator in people of a high level of mental health is "more efficient perception of reality and more comfortable relations with it."

I am excited about control theory precisely because the model offers "a more efficient perception of reality"—the way the world works, and we with it. In quoting Maslow, I am not trying to adhere religiously to any standard, rigorously defined control-theory terminology, if there is one. It seems to me that in addressing a larger public, I must find a way to use terms they relate to. So far in my attempts, I try to use as many synonyms and analogies as I can find.

Rick Marken: Dag, my comments about reality therapy and Glasser were not directed at you at all—nor were they meant as a criticism of your work. It just jogged a thought in my mind which I wanted to make public—about Glasser's interest in control theory. I've wondered why his therapy is called reality therapy if he is such a fan of control theory. I was asking the Net at large; there are a number of therapists out there, and some are familiar with reality therapy.

Again, I request info on this topic from therapists. I really am curious about it.

Dag Forssell: Rick, now that I understand your comment as a question, I shall attempt an answer.

Very briefly, Bill Glasser used/developed reality therapy 30 years back. His book by the same name is still available in your local book store. He developed an institute and a large following, numbering in the thousands. He has written many books which show his deep interest in matters human (see *Positive Addiction* and *The Identity Society).* He was told about our 1973 "bible" and attempted to write a version more accessible to the public.

I was fascinated by Glasser's *Stations of the Mind,* but then, I am an engineer. The book does a credible job, as far as I remember. He gives proper credit to Bill Powers. Many of Glasser's senior faculty still go by that book, which is why some are in the Control Systems Group. The book was probably not a hit with the public, and apparently not with most of Bill Powers' followers. Glasser developed a four-color chart to teach by, which is simplified but not bad.

Clearly, reality therapy came first. Control theory failed to support it as Glasser anticipated, since he could not teach it in a way his audience accepted. Problems of organizational control might have contributed to the break with Powers.

It seems to me that Bill Glasser is smart and has made contributions in many ways. Reality therapy is his baby and his dominant systems concept. It comes first. Glasser's book *Control Theory* provides the following definition of control theory: control theory contends that every behavior is a person's best attempt to meet his needs. Perception went out the window because it was confusing to his audience.

This is a quick sketch of my perceptions on this. We all have different contents in our systems concepts. Glasser's priorities are different from ours. Still, he has brought a number of control-theory faithfuls to our fold through his promotion efforts. I am glad that I am one of them.

Rick Marken: Thanks for the thumbnail sketch of the history of reality therapy and its relationship to the Control Systems Group. Actually, I am fairly familiar with that history. I went to a Bill Glasser show in about 1981 when I was in Minnesota. When I found out he was interested in Powers, I went right up to Bemidji or wherever he was. I even had lunch with him in the regal dining room—he invited me in when I told him I had been working on control theory for a couple of years already. He struck me as a consummately self-absorbed individual, not at all intellectually interesting. I still don't really understand the basis for the rift (if that's what it was) between him and Powers, though I would guess that it had much more to do with Glasser's rather shallow grasp of Powers' model than with any conflict over control of any organization (the very notion of Bill Powers trying to control some organization is pretty silly, given what Bill Powers is like).

Anyway, I still don't understand what "reality" therapy has to do with control of perception. I kind of ragged on Ed Ford for not explaining to my satisfaction why it is important to realize that people are controlling perceptions, but I think I react so strongly because it is so important to me that it be made clear. *Nobody* has direct access to reality. We control only representations of what reality might be. To the extent people can see agreement regarding what they perceive, we tend to call those perceptions reality—but they ain't. They are just (somewhat) shared perceptions. Indeed, I think using the term "reality" in a therapeutic situation could actually be dangerous—giving the therapee the impression that there is a right way to perceive things. My pedagogical point is that, when explaining control, just leave out the term perception *at first.* I control the letters on the screen, the position of my hands, etc., etc. Once a person understands that there are variables "out there" which are controlled (many different letters could be typed; many different hand positions are possible) and that these variables are brought to reference levels though the action of lower-order acts which could also counter the effects of disturbances to these variables, the person understands the phenomenon of control. Then you can explain that it is perceptual aspects of experience which are being controlled, and that there are, therefore, different perceptual aspects of the same experience which can be controlled.

Again I ask—really, just out of curiosity, not hostility—why would the person who developed "reality therapy" see perceptual control theory as something which would support his theory? And what is "reality therapy," anyway?

Ed Ford: For 14 years, I was a faculty and certifying member of the Institute for Reality Therapy. I taught at every Intensive and Certifying Week in L.A. (where most were held) for over 10 years. I was very close to Glasser. I left the Institute in 1983. Glasser once remarked (somewhat in jest) that he would have called what he did *The* Therapy, but he might have run into problems with others. He called it reality therapy because it best described what he was trying to do, which was to get people to deal with the reality of their present life. It was the most efficient therapy I knew at the time.

When I was introduced to control theory by Glasser, some of what he said didn't make sense (such as don't deal with perceptions, leave that to the theorists). At the 1989 convention mentioned below, Glasser said, when talking about perceptions, "they say it is a hierarchy and you always start out with this one and then get to this one and this one. I don't recognize 10 [levels]. I don't get involved with it. In terms of them [CSGers], it's a fundamental difference." Thus his total disregard for the hierarchy of perceptions (which he wrote about in *Stations of the Mind,* then obviously abandoned). I then left him and became a pupil of Bill Powers. When, after many years, I had finally begun to understand what control theory was all about, I realized that we control perceptions. Glasser never has.

I heard a tape recording of a workshop on what others are thinking and saying about control theory, given by Diane Gossen at the Institute's 1989 convention, which Bill Glasser attended. During the presentation, Glasser kept making comments and corrections to what Diane was saying. His degree of understanding of control theory was very evident.

When someone at the conference asked Glasser when does output become input, he replied, "the only way that the behaving organism becomes aware of the behavior is through its ability to perceive, which is input .... You can go through all kinds of outputs all the time, but what they [CSGers] are saying is that the only time you become aware of it [Glasser's idea of behavior] is through input." Glasser never has understood the concept of the controlled variable, and that we control for input. He sees "control of perception" to mean that when we perceive what we are doing, we are controlling our perception of our behavior. This is how he understands the title of Bill's book, *Behavior: The Control Of Perception.* He has never gotten away from the fact that we control our behavior. Remember, behavior to him is output, behavior to us is the entire system in operation.

He doesn't understand the levels of perception. Glasser said, "the reason I got rid of the levels of perception is when I started to teach you could adjust the levels, and I don't think you can. I think it is all the way through the top. You adjust to what the ultimate picture is... that's what drives you. If you think you can dissect your behavior... I think that is absolutely impossible." For Glasser, reference conditions are called "pictures in your head." The picture for him is the entire hierarchy of perception, and that is all you can control. Obviously, the entire system is engaged in the operation, but he doesn't believe you can be aware at any one level. Another major problem is that he uses "picture in your head" interchangeably as both perception and reference level. He doesn't see the comparison going on inside the head (between perception and reference condition), but rather between the picture in the head and what he calls behavior (and what we call our actions). At one point, he uses the picture-in-your-head concept as building a perception; at the next, as something you want or a reference condition. The same word is used for two entirely different concepts. The bottom line is that Glasser has never, never gotten away, from controlling output. For him, the comparison is between what we want (which he calls the picture in your head) and what we are presently doing to get what we want (our actions or what we are doing which we call output).

Other areas of misunderstanding by Glasser: He says his idea of needs is what CSGers would call disturbances. Obviously, he doesn't understand disturbance, because he doesn't understand the concept of controlled variable. Glasser sees the reasons for disturbances occurring as the basic needs. And there you have another major problem, the concept of needs. There are basic needs such as the need for water, food, etc. Where it gets tricky is when you get to such needs which Glasser identifies as Power, Fun, Freedom, Belonging. Our genetic system sets the limits on basic needs. But when it comes to those areas through which all of us strive to find satisfaction, they can be seen quite differently by each of us. I really struggled with this idea (with a great deal of help from Bill Powers) in Chapter 7 of my book, *Freedom from Stress.* It seems to me that we set the limits and parameters of satisfaction within our own hierarchy, especially at the higher levels. This setting of limits is really our individual mark or standard for areas of importance to us, what Glasser would call needs. I think Glasser's higher-order "needs" say more about him than anything else. People define their own internal goals and areas of satisfaction, and from my daily reading of this Net, they surely vary a lot within our own Group.

Other areas, which I don't want to dwell on, include the following: He calls feelings (along with doing and thinking) behavior (remember, his definition of behavior is output). His understanding of reorganization is also very confused. He retitled his book *Take Effective Control of Your Life,* calling it *Control Theory.* The sad thing is that he has taken the name control theory and assumed total control of what it means.

Bill Glasser taught me more than anyone else a lot of great techniques for counselling and dealing with others. Unfortunately, when he was exposed to control theory, he changed control theory to suit his own perception of the world, and to suit reality therapy. Over the past eight or 10 years, as I have been learning control theory, I have tried to take my ideas of reality therapy and other ideas in counselling, and adjust them to the new and very different world of control theory. Control theory has opened a whole new understanding of the world to me, and thanks to Bill Powers, Tom, Rick, and the rest of you, I have been able to use control theory to learn and grow as a counselor, father, teacher, husband, and all the other hats that I wear.

The basic tenets of reality therapy are the steps (get involved, ask what do you want?, what are you doing?, and is what you are doing getting you what you want?; then get a plan and commitment). It also says that whatever we are doing is our attempt to satisfy our needs. I have been able to use reality therapy as a jumping board from which to develop a control-theory therapy.

David Goldstein: Many years ago, I attended a workshop which Glasser gave about the time he started to publish books about control theory. I asked him the very question which Rick raises. From his answer, I received the strong impression that as a result of studying control theory, he revised his attitude from emphasizing a more objective view to a more subjective view of perception. Ed Ford and others who started out in the reality therapy camp could probably tell us more about it. It seemed that he found the questioning a little discomforting. I will leave it to Ed and others in the Control Systems Group to tell us how reality therapy compares to control-theory therapy. I am writing to tell about "experiential therapy."

I have recently read a book by Alvin R Maher, called *Experiential Psychotherapy: Basic Practices.* I would highly recommend this book to other control-theory clinicians. While the theory of human beings behind experiential therapy is not control theory, but rather a form of existential-humanistic theory, I think that much of the practice is very consistent with control-theory therapy. For example, the method of levels is there! And an effort to describe a particular person's control-system hierarchy is there!

The meaning of the word experience in experiential psychotherapy is very much like the meaning of perception in control theory. I am going to focus on the ways in which experiential psychotherapy can contribute to the practice of control-theory therapy, rather than vice versa. I do see control theory as making a contribution to experiential psychotherapy in many ways, but I will save that for another time and place.

To start with, in experiential psychotherapy, the selection of topics is based on the intensity of experience which the patient experiences when attending to the topic. From the point of view of control theory, this makes good sense. Life areas in which a patient experiences a lot of or only a little bit of satisfaction are likely to be ones about which a person has strong experiences.

While they seem to select topics in similar way, experiential psychotherapy and control-theory therapy differ in the way a topic is pursued. It is here, I believe, that experiential psychotherapy can teach us. The therapist attempts to share the patient's experience as much as possible, including the bodily experiences which go with discussing a topic. Control-theory clinicians do not do this, as far as I know; they do not try to generate the patient's experiences within themselves.

The therapist attempts to have the same experiences as the patient is describing. This is called the "working level" of experiencing for the therapist. When this is not occurring, the therapist and patient stop and go back to the point at which the therapist lost touch with the patient. The sameness of experiences includes bodily experiences!

Maher provides some specific methods for helping a person achieve a strong experience, for helping a person become aware of the higher-level experiences behind the one being discussed, for helping a person be/behave in a way consistent with the higher-level experience, and for helping a person extend the changes in being/behaving to the patient's everyday life. These specific methods are very helpful, are based on years of clinical practice, and I see no reason for not using them.

In closing, I wish to emphasize that experiential psychotherapy makes me rethink the way I am using feeling/mood experiences in the control-theory therapy which I do. I no longer simply try to intellectually figure out (imagine) what the patient's blocked desire is, which is generating the feelings/moods. I try to experience the feelings/ moods and the topic being described by the patient. I then use the experiences which are occurring within me, which are not explicitly described by the patient, to experientially figure out what the patient's higher level perceptions are. In short, experiential psychotherapy is like control-theory therapy, but with more feeling for both the patient and the therapist.

Now, back to reality therapy. While I am not a trained reality therapist, I have watched Ed Ford, Diane Gossen, and Perry Good work at the CSG meetings. And I have read their books. It is my impression that feelings/moods, experiences, and expressions also play a secondary role in the therapy which they do. Reality therapy plugs into the last step in experiential psychotherapy.

Rick Marken: Thanks to Ed and David for the info on reality therapy. I think I get why Glasser got interested—just a poor choice of words, that "reality" thing.

Bill Powers: David, I presume that you remember a discussion in which we talked about devising "qualitative models" for kinds of behavior which are difficult to handle in a purely quantitative way on a computer. Your description of Maher's *Experiential Psychotherapy* strongly suggests that it is a way of doing exactly what we were talking about. The "computer" in which you run your simulation of the other person is, of course, yourself. This living computer already contains the capacity to carry out, in imagination, all of the functions of a human being (oddly enough) at exactly those levels of functioning which actually exist. There is no programming problem—we don't have to figure out how relationship-perception works, or program-perception, or system-concept-perception. The computer is guaranteed to be able to run any process at any level which is required. It is also guaranteed to contain exactly the levels which are required, not skipping any and not adding any which don't belong in a model of an adult human being.

As I read your description, I was reminded of the problem of listening to someone who is giving you directions. When I hear the directions, I try to imagine the actual trip being described. Good directions give you a picture vivid enough so you don't have to write anything down—when you actually follow them, it's as if you've already been there. Poor directions, on the other hand, are full of skips and jumps, private associations and incidental anecdotes; they convey a shifting point of view, sometimes from the viewpoint of the one taking the journey, and sometimes as if from an aerial map or the position of a by-stander.

When you try to follow poor directions in your imagination, you get a picture of a very confused mind. You don't, of course, actually sense the other person's confusion. But by trying to imagine following the directions, you become confused yourself. That is, when you try to run the model the other person is describing, it leads you to see gaps and contradictions and other problems which leave you confused and, by implication, indicate at least a similar kind of confusion in the other person.

I think that this method can be refined by a control theorist into something even more workable than it already is, and also that it can tell us a lot about the role of language in the control-theory model. Language, in the broadest sense of communicating through manipulation of perceptions, is the medium through which one person tries to convey his or her experiences to another. It undoubtedly has limitations—there might be inherent difficulties in trying to communicate principles and system concepts by any means but demonstration, and in trying to communicate very low-level perceptions, such as the way a face looks or the way ice cream tastes. There are problems inherent in private meanings of words. But in an intimate and protracted relationship with one other person, a therapist should be able to cross-check the meanings and put together a quite reasonable model of the other person, through imagination.

This means that the therapist must become an utterly flexible general-purpose simulation device without cultural biases and with no blind spots—a selfless person. At least during the process of therapy. You would not want to simulate every person as if he or she were, for example, a middle-class Jew or an eccentric engineer. You have to allow the properties of the other person, as nearly as they can be communicated, to enter into yourself and to operate as if they were yourself. It seems to me that doing this would amount to a discipline at least as rigorous as that which the Zen masters demand, at least as deep as the analysis which psychoanalysts are required to go through before they are considered ready to treat other people's problems.