This reproduction of *Closed Loop* was created by Dag Forssell in 2001. Addresses and phone numbers have not been updated. Most are obsolete.

Posted at www.pctresources.com

Proofread as of March 6, 2001

Closed Loop
Threads from CSGNet

As announced in the last CSG newsletter, CSGNet, an electronic mail network for individuals with control theory interests, was begun in August 1990. CSGNet now has about 40 participants in the U.S.A., Canada, Europe, and Australia, and it continues to grow steadily.

Since its beginning, CSGNet has been a remarkably active forum for the discussion of control theory. CSGNet has turned out to be an exciting and convenient medium for sharing ideas, asking questions, and learning more about control theory, its implications, and its problems. Among the more active CSGNet participants are Bill Powers, Gary Cziko, Rick Marken, Wayne Hershberger, Tom Bourbon, Chuck Tucker, David McCord, Dennis Delprato, and Hugh Petrie.

A serious shortcoming is that to date there are no clinical participants. This is most likely because most CSG clinicians are not affiliated with university or research institutions having access to either the Internet or the Bitnet electronic mail networks. Nonetheless, at least one commercial computer communications service, CompuServe, offers access to Internet (and therefore to CSGNet) for its subscribers. This means that independent researchers and clinicians who do not have institutional access to Internet or Bitnet can still participate in CSGNet. They just need a computer, modem, telecommunications software, telephone line, and money to pay for the connect time.

As this is written (January 1991), CompuServe's connect time charges are $6.00 per hour for 300-baud service and $12.50 for 1200- and 2400-baud service (call toll-free 1-800-848-8990 for up-to-date information). To obtain access via a local telephone number in most American and Canadian cities, Telenet is probably the least expensive telecommunications service to link to CompuServe. Telenet currently charges $12.00 per hour during prime time but only $2.00 during non-prime time (evenings and weekends). To make the most use of CSGNet at minimum cost via a commercial service, participants should connect to the service only for uploading and downloading mail. In this way, messages can be composed and read off-line. Two lower-cost services that do not now have access to Internet but might have access by the time you read this are GEnie and Prodigy. GEnie currently charges only $4.50 per month for unlimited access to its basic services, including electronic mail. There are no sign-up or connect time charges for participation on CSGNet itself.

CSGNet's Bitnet address is "CSG-L@UIUCVMD " (use no quotes in this and the following addresses); "CSG-L@VMD.CSO.UIUC.

EDU " is the address for Internet. The messages sent to CSGNet via these addresses will be forwarded automatically to all participants. Use the address ">INTERNET:CSG-L@VMD.CSO.UIUC.VMD " to reach CSGNet via CompuServe. To become a CSGNet participant, initially send a note to the network manager, Gary Cziko, at "CZIKO@UIUCVMD " (Bitnet) or at "G-CZIKO@UIUC.EDU " (Internet).

Gary Cziko
217-333-4382

Last October, when CSGNet was just taking off, Bill Powers and Tom Bourbon and I talked about publishing a sort of patchwork newsletter stitched together from the Net's conversational "threads." I agreed to consider the feasibility of such an undertaking and, if interest seemed high, to start it off with the understanding that Tom and others would provide assistance or take over as they were able. I've been impressed by the highly creative and substantive dialogue on the Net, and I suppose that its quality will continue to flourish as more participate in the discussions. At least to date, there has been plenty of material worthy of preserving and disseminating in a "digest," which would also allow Net participants—at the behest of the editor—to clarify and expand their comments in light of reactions to them by other Netters.

The question, of course, is whether anybody else is excited by the possibility of a CSGNet digest. The following "threads" from the Net will give an idea of what can be expected in a digest, except that for this "sample issue" of the digest, I didn't ask participants to elaborate on their original statements. Please let me know if you think the project is worth pursuing to the extent of at least one full-size issue. Would you pay $10.00 per year for two issues of *Closed Loop*, each about 100 pages (like these) long? Would any of your non-CSG colleagues and/or local libraries be willing to pay $20.00 per year? I'd appreciate any and all comments and suggestions.

Greg Williams
606-332-7606

The Uses of Control Theory

Rick Marken: Many people have the idea that the true test of the value of a theory is whether it is "useful." This seems to be particularly true in the field of psychology. One unquestionable reason for the popularity of behaviorism is its apparent usefulness: it tells you how to cure "behavior problems;" raise children, manage people, etc. I think a case can be made for the proposition that cognitive psychology (and its vari-

ants) really came into its own when it learned how to sell itself as "useful." Thus, the popularity of Al (with helpful expert systems), human-computer interface engineering (my own field), neural nets, fuzzy logic (the Japanese use it in washing machines!?!), etc., all of which are related to cognitive psychological theorizing. Even Freudian, Jungian, and other "clinical" theories are popular because they promise to show you the source of your own problems: they claim to be useful.

I am often asked, when I present control-theory ideas, "so, how can I use this; what will the theory buy me?" I don't think that I happen to be running into an unusually utilitarian group of people. I think all people look at ideas in terms of what those ideas can do for them: after all, people want to be able to control things better; they are control systems. The success of science in general (and of scientific theories in particular) is typically presented in terms of "look what science has made it possible for us to do (control)." Science is seen as the handmaid of control; not as a window on understanding.

I submit that people's interest in "usefulness" puts control theory at a huge disadvantage in the public eye. Things that are useful help us control. But control works best when practiced on objects that are not themselves trying to control. Control theory tells us that people are trying to control. Unfortunately, people have the nasty habit of mistaking "other people" for the kind of objects that can be controlled. One of the main goals of control theory (as I see it) is to teach people that other people are not that kind of object (the kind that is more familiar to physicists). In fact, control theory suggests that efforts to treat people as though they were controllable objects are likely to lead to conflict rather than success. When there is conflict, there is no control on either side.

I don't think that the message of control theory is "just leave people alone and everything will be all right." But the message is definitely not "if you understand control theory you can get people to behave just the way you want." Many of the people who have asked me about the uses of control theory have definite goals regarding how they want people to behave. These people tend to ignore a theory if it doesn't say "in order to get behavior Y you do behavior X." It is difficult to convince them that, in the long run, they will be able to achieve their goals more successfully if they are more selective about what they try to control (non-living systems) and what they try to cooperate with (living systems).

So, what do you think? What is the use of control theory? How would you communicate its usefulness to, say, an experimental psychologist, the manager of a business, a plain old ordinary person?

Gary Cziko: I think control theory can be very useful for education, management, and clinicians (as demonstrated at our meeting), but

there is also a scary side as well.

If control theory tells us that attempts to control other people using "peaceful coercion" ultimately lead to conflict and violence, then why not start with conflict and violence from the beginning? Saddam is now controlling the oil production of Kuwait quite successfully by using force.

Rick Marken: I don't think the control theory message is that "peaceful coercion" will necessarily lead to conflict. In fact, peaceful coercion could be quite successful. Actually, it seems to work all the time. I want to eat and this company is willing to give me money so that I can. I understand that my role is to "work" for them—where "work" can mean spending my time doing something that I prefer to do less than other things. I'm willing to make this exchange—the company "controls" what I do, and I control the amount of money I get. It works because, so far, we are both willing to accept a little error—I don't get nearly as much money as I want, and they probably don't get all the work they would like to get out of me. But we're both happy.

Control theory just says that when you deal with a person, you are dealing with a control system. The result of that "dealing" depends on how you deal with the control system and what the control system's current configuration is. But it is true that if you try to control the control system "arbitrarily" (that is, without taking into account its purposes), there is a good chance of conflict. For example, if the company decides that it will only pay me if I work in a certain way, and if it's the only company in town and I have no alternative means of getting money, then there are likely to be problems if, for some reason, I don't want to work in that particular way. If the company tries to control me —meaning it will only accept seeing a particular kind of behavior on my part—and if that behavior is something I just don't want to do, then there is conflict.

Most people deal with other people as people—they act as though they understand that the other person is a control system and they show respect. You get into problems with very "purposeful" people who have to have people behaving in just a certain way—no attempts at cooperation. These people treat people as objects. When I control a hammer, I want it to move exactly as I want it to move. I don't want to compromise and say, "well, if you want to land a few millimeters closer to my thumb then it's OK with me—I understand that you have needs too." I don't say that because the hammer has no needs or wants, and I can control it perfectly—we never have conflicts. But if I act the same way with my daughter, son, or wife, I am probably looking at significant conflict.

You brought up our current crisis with Hussein. How would you

analyze the situation from a control-theory perspective? Obviously, Saddam is an example of the kind of person I described above as "purposeful." He clearly wants something, and he is willing *to* engage in conflict in order to get it. I argue that conflict can never be a good solution, even for the victor, since strong control systems will prevail over weak ones in a conflict. Conflicts are most interesting and obviously debilitating when both parties are about of equal strength (or skill or whatever). But even the winner of the conflict is a loser (in the long run). It is very seductive—winning a conflict looks like successful control by the person who does win. But I argue that it is a fool's paradise. The winner then imagines that control can always be achieved by force (not true), and the loser never really goes away.

I admit that there are many instances where the havoc being wreaked by a control system is so bad for other control systems that there seems no option other than forcible conflict (Hitler comes to mind, slave owners, and possibly Saddam). But can't we think of ways to avoid getting into these situations? I just can't believe that there are that many "evil" control systems running around.

David McCord: Rick, your interesting remarks suggested to me a potentially very useful aspect of control theory—conflict resolution. Conflict situations are often those in which two parties are controlling the same input quantity around different, incompatible reference levels. From a control-theory perspective, though, we know that those reference signals are merely the means to ends, outputs of control loops one level higher. Conflict resolution typically involves "going up a level" in order to identify higher-level goals of each party that are not fundamentally incompatible. While this technique is included in many different approaches *to* conflict management, control theory provides a unique understanding of why the technique works.

Chuck Tucker: I believe that the major argument for the usefulness of cybernetic control theory (or what I call Sociocybemetics) is that it is a model of how a system and process work. This is the point that we have made over and over again in our meetings—the model tells you and everyone how living systems both individually and collectively work—how they do what they do-how to fix something when it goes wrong—how to make it possible for a system to destroy itself (positive feedback)—how to suggest a system solve problems-how problems can be located—and much more. This is basically the argument for the type of model we use and it differs drastically from the types of models (theories) that are used by almost everyone in the life, social, and behavioral so-called sciences. Now perhaps we need to catalog or collect illustrations, examples, and stories about how the model has worked, so

we can have them handy to present to persons with whom we interact. I suspect that this network would be a good place to begin our list of *working examples of CCT. How* about it, mates???

"Revolutionary" Control Theory?

Chung-Chih Chen: I have read "A Manifesto for Control Theorists" by Powers. It is really very interesting. I like the idea of being a revolutionary. That is always what I want to be. But it seems to me that it's very apparent that a living system can be regarded as a (feedback) control system used in engineering. So I am very surprised that the manifesto claimed that it is a new idea for life science. I wonder why life scientists didn't discover it before.

Rick Marken: What is new, I think, is that the control of perception (which is what feedback control means in organisms) is the fundamental organizing principle of living systems. It is the fundamental organizing principle because what living systems do, at all levels of organization, from the cell to the organismic level, is carry out purposes—i.e., they control. It is the fact that organisms control, rather than what they control, that is of central importance to control theorists. Control theorists are more impressed by the fact *that* organisms control than by *what* they control. It is just as amazing that a cat controls the texture of the food it eats as it is that a person controls the network of contingencies that produce checkmate in chess. It is the organizing principle that is revolutionary: behavior is the control of perception.

AI types seem to be more impressed by the kinds of complex variables that people can control than they are by the phenomenon of control itself. This is certainly understandable. I'd rather watch my kid play chess than watch my cat chew. It is the content of control, rather than the organizing principle, that interests AI and cognitive science types, in my opinion. But Al types certainly know about control theory and some have a pretty good feel for what it is about. I was just looking over Minsky's *Society of Mind* book. He has a couple of chapters on "difference engines" which reflect a definite understanding of the purposeful nature of their behavior. (A difference engine is just a feedback control system.) He definitely understands that these systems produce goal results in the face of disturbance. But he doesn't really grasp the idea that this means that they are controlling perception, not "output." So near, yet so far.

Ultimately, AI and cognitive science seem to have concluded that control theory is just a subcomponent of a more overwhelming model of

human nature. I think if you look carefully you will find that this overwhelming model is some form of external causation—where "external" could mean in the environment or in the brain / nervous system. Just like the behaviorists, AI people often get very close to the underlying principle of control (purpose) and then go off and do something else instead. Still, much of the AI / cognitive work is relevant to control theory. I see it as explorations of some program-level perceptions that people control and how they might control them. They also are more explicitly concerned with control of self-produced perceptions (those not produced via the external loop through the environment), such as memories and imaginings.

So, finally, the control revolution is really based on taking purpose seriously and understanding that purpose must be organized around the control of perception. For research purposes, this means that a large part of understanding the human mind must involve learning the nature of the perceptual variables that it controls.

Wayne Hershberger: Chung-Chih Chen, welcome! I understand your incredulity. I still do not understand how psychologists can fail to recognize the fact that animals control their environments, to the degree that they are able. Indeed, we are all puzzled; read the introduction of William T. Powers (1978) Quantitative analysis of purposive systems: Some spadework at the foundations of scientific psychology. *Psychological Review 85,* 417-435.

Tom Bourbon: When Chung-Chih Chen expressed surprise that the life sciences don't embrace control theory, Rick replied that they are close to the model, but are not quite there. I'm not sure I agree, at least not entirely. It seems to depend on which sources you read. If you look at accounts in physiology and in "neuroscience" of the control of movement via skeletal muscle, then there is little doubt that few life scientists appeal to control theory as an explanation, and that many of them reject the control model.

But the picture is quite different when the discussion shifts to internal variables. There, for several years, many physiologists have used a fairly good control-system model. Not the old, rather static models of "homeostasis," but models in which the "set point" (our "reference signal") is compared to a negative feedback signal from sensors that detect the present state of a controlled variable. And the present state of the controlled variable is a function of the output of the system (they now recognize that the external variable, not the output function, is important) plus the effects of disturbances of all sort. If you want a good representative text, try *Human Physiology,* R.F. Schmid and G. Thews (Eds), Springer Verlag (1983). There are many more. This version of a control

process is so widespread that most authors do not even cite a source—it seems to be taken for granted.

The biggest differences I see between their models and ours are these: they still refer to a comparator as a controller; the error signal is still called a command signal; and the perceptual signal is their negative feedback signal. And they do not yet realize that the perceptual signal is the variable the system really controls. Of course, we don't help the situation very much with our terminology—calling the external variable the "controlled variable," then chastising people when they do not realize that the system controls its perceptual signal, is not terribly fair on our part.

As for cognitive models...! If there were any remaining doubts that they reduce to S-R models in I-O model clothes, those doubts are over. Read "What connectionist models learn: Learning and representation in connectionist networks," S.J. Hanson & D.J. Burr, *Behavioral and Brain Sciences* 13(3),1990, 471-518. On page 473 is a re-creation of Egon Brunswick's old "lens model" in which many environmental "inputs" converge on, and are "focused by," a lens (now called "unit processing"), then there emerge many expanding outputs. The inputs are now called "fan in;" and the outputs, "fan out." I'm not sure the model explains anything more than Brunswick's did.

More important, the authors dearly identify the goals of con- nectionist modeling, as they see them: to show how the "hidden layers" in the model allow it to match outputs to inputs. There it is, clear as day, the thing we have known all along, but were criticized for saying: most "cognitive" models reduce to stimulus-response models by another name. The implications of this fact are great, given that cognitive-neuroscientific theorists declare behaviorism "dead," and their models both superior and ascendant. And a majority of them view control models as just another version of cybernetic feedback models, able to account for only a portion of "mere" sensory-motor coordination, if even that. (See especially their remarks on p. 472, right-hand column, and p. 481, right-hand column.)

Rick Marken: Tom, I think we agree more than you think. I do think that the life sciences are often close to control theory (in my perception of closeness) but, in science, a miss, even a near miss, is a mile. The reason they are close (in my perception) is because a stimulus-response model can look an awful lot like a control model. It can even behave like one! And, as you correctly point out, the model that the life sciences are ultimately trying to defend is some version of a stimulus-response model.

A stimulus-response (or response-selection) model works when you *define the stimulus in a way that implicitly includes the reference condition.*

The stimulus-response model works because behavior is occurring in a closed loop. So the model can be called a stimulus-response model, but it is really a control model with the reference signal implicitly set to zero. An excellent example of this same thing can be found in some work on computer animation that I have stumbled across. Here are some references for those who are interested: J. Williams and R. Skinner (1990) Motion Control: A notion for interactive behavioral animation control. *IEEE Computer Graphics and Animation,* May, 14-22; V. Braitenberg (1984) *Vehicles,* MIT Press; C.W. Reynolds (1987) Flocks, herds and schools. *Computer Graphics (Proceedings of SIGGRAPH),* 21, 25-34. These folks have built little control systems that follow things or move to targets on the screen. But they don't think of them as control systems; they have sensors and effectors, so they "must be" stimulus-response devices. The devices exhibit some pretty impressive, goal-seeking activity. These researchers are sure that they are S-R devices with no inner purposes. But they are actually control systems. The sensor input does affect the effector output, but the effector output also affects the sensor input; there is a closed loop. The loop is stable because there is 1) negative feedback, because they have set up the S-R rule so that the output nulls the input, and 2) proper dynamics; there is slowing of the output effects of the sort that we use when we write our models of control. That is, the output at time t is proportional to the integral of the stimulus over time.

These "stimulus-response" devices are really control systems. They will reach their targets even when there is disturbance. But they illustrate what I mean by "close, but no cigar." These people are building control systems and watching them behave purposefully. But the researchers don't see this because they are guided by the unseen principle that behavior must be guided by external events.

One thing that might be fun is to build some of these simple organisms, but put in an explicit reference signal. This should be a variable reference signal, and, for now, it could just vary slowly and randomly. Now we have an organism that is still "S-R" in the sense that these researchers imagine, but one which is always clanging the definition of the stimulus on its own. The random changes in the reference produce "spontaneous" behavior that cannot be controlled by an external observer. But it is possible to demonstrate that the behavior is still purposeful and organized (nonrandom) by applying disturbances and seeing that they are resisted.

Tom Bourbon: Rick, as for how close the life sciences might be to an understanding of control, look again at the reference I cited as an example. There are *many* similar examples. These people are *not* talking about motor control. Instead, they are describing the control processes

for internal variables. In that field, the understanding has progressed dramatically since only a few years ago–so much so that we risk alienating a very large community when we say, flatly, that the concept of control is not understood in the life sciences. The reference signals (a.k.a. set points) are explicit, not implicit; the output is not the object of control, rather, there are *clearly identified* controlled variables (external to the control system); disturbances affect the controlled variables; and so on. Obviously, these are not the people who reviewed our manuscripts!

Look at the Schmid and Thews reference or at one of Mountcastle's more recent editions of *Medical Physiology*. What you will see in no way resembles the literature on motor control, or most of the literature on "cognitive neuroscience." I think you will be pleased: it is science, not seance.

Chung-Chih Chen: Thanks for all comments on my surprise. I am looking for the suggested papers and studying the feedback control system from the beginning. I am still not sure if control theory is a revolution. I will tell you when I understand better.