

The Tank that Filled Itself

William T. Powers

Copyright 2010 by William T. Powers.

Reprinted with permission.

This paper was presented by William T. Powers at the Control Systems Group Conference in Manchester, UK in July 2010. It is posted on www.nickols.us with his permission.

Perceptual control theory (PCT) is a cross between biology and engineering. Its shallowest roots are in *control theory*, developed during the 1930s and 1940s by electronics engineers; and *homeostasis*, physiologist Claude Bernard's idea from the mid-1800s. Walter B. Cannon carried Bernard's idea further in the 1930s, and a decade later Cannon's ex-student Arturo Rosenbleuth, aware of new engineering developments, told Norbert Wiener about the resemblance of control systems to behavioral systems in human beings and animals. That was the start of cybernetics, the science of "steersmanship." PCT was probably conceived when I learned something about control systems first as a navy electronics technician and then as a student physicist. It was brought to life in 1952 when I read Wiener's 1948 book.

But the tap root of PCT goes far deeper than the strata in which we find the control engineers of the 1930s, or Wiener, Rosenbleuth, Cannon, and Bernard, or me as a young man of 26. It burrows through layers of engineering developments in which we find 19th-Century control systems for steering steamships, down through Watts's 18th-Century flyball governor for steam engines, through the wind-driven grain-mill speed regulators ("lift-tenters") immediately beneath, through medieval temperature controls for furnaces, down through Arabian water clocks, all the way to a Greek inventor named Ktesibios, a student of Archimedes, a contemporary of Euclid, and possibly head of the Museum of Alexandria in Egypt before the great library was burned. Ktesibios was interested in water clocks.

The road not taken: the first recorded negative feedback control system

Water clocks, in ancient Egypt, had a small tank that held water used to fill, very slowly, a larger reservoir, raising a float and moving the time pointer. To keep the smaller tank filled precisely to the right level, so the clock would keep proper time, ancient Egyptians had either to keep the tank filling fast enough that it continually overflowed, or to assign a slave to replace a much smaller amount of water as it was used by the clock. In about 250 BC, Ktesibios thought of an automatic device that would prevent wasting water and making a mess. His regulator would replace the slave, using a float in the small tank to measure the water level and a link from the float to operate a valve that would let more water in when the water level dropped. Now the tank could keep itself full without human help.

With a little poetic license, we can imagine a pre-Ktesibios water clock outside in an ancient courtyard, but equipped with the Ktesibios regulator. The regulator would respond to all kinds of disturbances in just the right way. If a flock of birds took a drink from the small tank, or if water slowly evaporated on a windy hot day, the float would fall slightly and the valve would open a bit more to compensate for the transient or continuing loss. If a rainstorm overfilled the tank or somebody tossed a stone into it, the valve would close until the clock's use of water lowered the water level, then open just enough to keep the tank at the former level. Everything this device did, the slave it replaced would have done.

It's important to notice that neither the slave nor the regulator had to know anything about why the water level varied. They didn't have to chase away thirsty birds or people throwing stones or anticipate hot dry winds. All either one had to do was sense and affect the very thing that was supposed to be controlled, the water level. The slave sensed it by looking; the machine sensed it with a float. If the water level went below the right level,

The Tank that Filled Itself

the regulator, human or mechanical, was internally connected so as to open the valve until the intended level was restored, then to adjust the valve to maintain that level. The valve was being opened or closed as a means of controlling the *sensed* or *perceived* water level, since that is all the slave or the machine knew about the actual water level. We can say that the actions were the means by which either the slave or the machine controlled a perception of water level based on the actual water level.

Here is a diagram of how the Ktesibios regulator worked, which is certainly in the running as a diagram of how the slave, if there was one, would have worked, too.

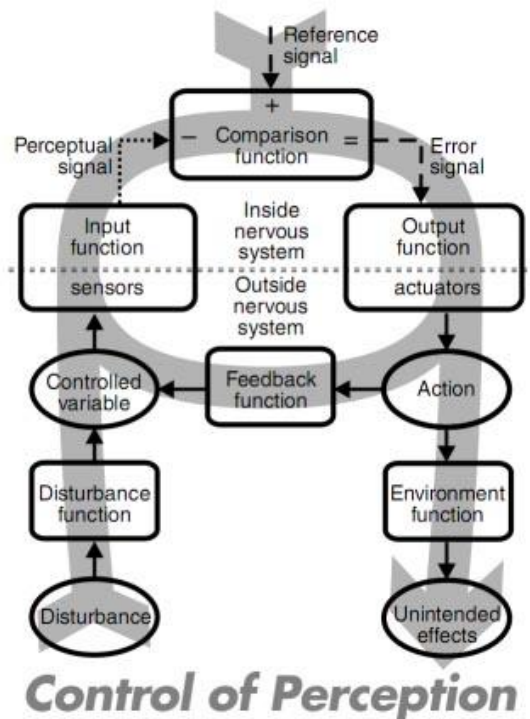


Figure 1: The basic diagram of a perceptual control system. By Dag Forssell, 2010.

The disturbance at lower left is a bird or a stone or a rainstorm or a leak in a pipe or any of a thousand phenomena in the environment of the clock. Each of those phenomena, to act as a disturbance, would have to be connected in some way to the water level in the small tank, the controlled variable in the diagram. The "disturbance function" represents the means by which each disturbance changes the water level -- removing water, displacing it to raise the level, or adding water.

The regulator proper is above the inside-outside boundary marked by double dashed lines. There is a sensor (vision or float) in an input function that converts the controlled variable into a perceptual signal, an internal representation of the water level. This signal is compared to a reference signal, which is an internal representation of the desired or

The Tank that Filled Itself

intended water level. In the artificial regulator, the reference signal is simply the water level at which the valve is just barely closed -- it can be changed by altering the link that connects the float to the valve. In the slave, it's a neural signal that biases the relationship between sensory input and output action, or more simply, it's a mark on the wall of the tank selected as a target for the water level.

Out of the comparison function comes a signal or an effect that is converted by the nervous system of the slave or the construction of the device into an output action. Opening or closing of the valve depends on whether the error signals too little or too much water, and the degree of opening depends on the size of the error. We are now back in the environment of the regulator. The action operates the valve, which in the diagram is called a feedback function, since it feeds an effect of the output action back to the input sensor. The action also has other effects, shown lower right, such as exciting the optic nerve of someone else watching the action, but those effects (which the observer experiences as the behavior of the regulator) have nothing important to do with this particular regulator. The effect of action that matters is the effect on the controlled variable, the same variable that the system is sensing. The water is maintained at the intended level, the "reference level," by purposive actions governed by a goal, the goal being a specification for the state of a perception that is to be created and maintained.

The road that was taken

Psychologists and neurologists and biologists were not stupid in the first third of the 20th Century. They were simply unlucky. They did not happen to think of the arrangement Ktesibios thought of, nor were they in contact with the engineering projects in which the ideas of Ktesibios were being elaborated upon at an increasing pace. The main result was that when they saw behavior that looked as if it were goal-directed by the behaving system itself, they ran into what looked like an impossibility. If they had lived in Ktesibios' time, they would have asked how the final water level, which did not yet exist, could reach back through time and alter the water intake so as to cause that specific level to come into being. Aristotle had spoken of "final causes," which included the way the idea of a chair could cause a carpenter to shape wood so as to build one. But in 1910 nobody believed in those primitive concepts any more. The scientific consensus was that this appearance had to be an illusion, and that when we understood the brain better we would see that the slave was not intentionally regulating the water level; he was simply responding to stimuli in such a way that the change in water flow happened to compensate for whatever had caused the change in water level. As to how Ktesibios' clever little trick worked, that would be a matter for engineers, not psychologists, to investigate. It couldn't have anything to do with how the slave really worked.

Behaviorists side-stepped this problem entirely early in the 20th Century. They decided in effect that we had to leave questions like this to future generations when we would know more about how the brain works. Until then, all that a scientist could do was to observe effects of the environment on an organism, record the behavior that followed environmental stimuli, and thus elucidate the laws of behavior. It seemed obvious that if there were no stimulus inputs, there would be no behavioral outputs; that became a matter of scientific

faith. Organisms could not initiate anything. Like any object made of matter, they could only respond to external forces and influences as their histories and their internal construction dictated. This was the line of thinking which, however reasonable, set behaviorism on the path to extinction.

Shortly after the 1930s when control engineering came into existence, psychologists who still wanted to explore the mind inside the organism, not just behavior, organized a new approach called cognitive psychology. Fighting the scoffing of behaviorists all the way, they tried an orderly approach to studying the internal organization behind behavior, if not in terms of mechanism then at least in terms of function. They started to make models, computer models into which they could put functions they assumed to underlie behavior, in an attempt to demonstrate artificial but intelligent behavior rather than just responses to stimuli. These psychologists knew, of course, about cybernetics and control systems because those were major topics in the 1940s and 1950s. But they didn't know enough about control systems, and tried to invent their own explanations of purposive or goal-driven behavior. Many cyberneticists joined them, but cyberneticists hadn't learned much about control systems either.

If these cognitive psychologists and the cyberneticists who joined them had lived in the time of Ktesibios, they would have explained the slave's behavior in a new way. Instead of seeing just stimuli and responses, they would have envisioned a complex set of brain functions at work. First there would be a goal, a desired state of affairs: the water is to be kept at a specific level in the small tank. The brain would have to detect incipient disturbances that could change the water level, and predict the amount by which they would raise or lower the level. Then given this number, the brain would calculate the way muscles would have to change joint angles of the limbs and fingers in order to cause the valve to open or close by the right amount and for the right length of time to replace the lost water or to let the excess water drain out into the main reservoir. With this plan of action completed, the brain would then, just as the effect of the disturbance arrived, issue the required neural signals which would operate the limbs and the hand to turn the valve one way or the other and then, if needed, back to normal.

Perhaps cognitive psychologists would not have accepted this explanation if they had ever seen Ktesibios's regulator in operation, where obviously none of those processes was happening yet the result was exactly what was needed. But they did offer similar explanations of other behavior which are still believed by a very large number of life scientists (excluding me).

I have delayed showing the system diagrams of behavioristic and cognitive models until now so we could see them side by side.

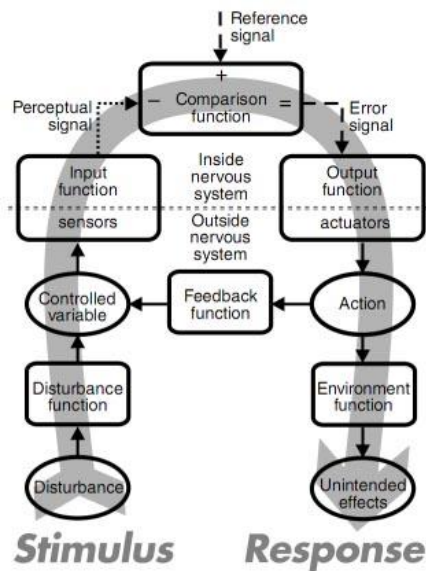


Figure 2: Behaviorism: Stimulus-Organism-Response.

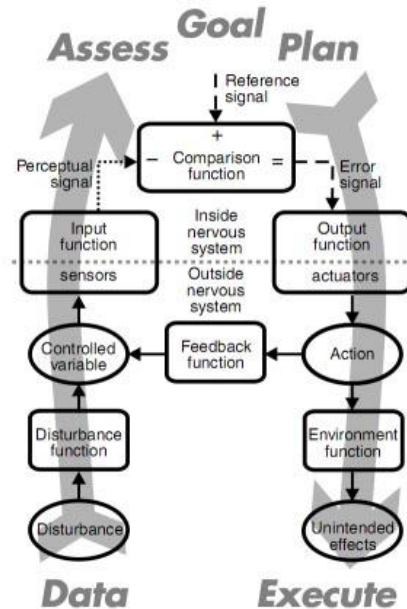


Figure 3: Cognitive Psychology: Data-Assess-Goal-Plan-Execute.

These models are quite similar, differing mainly in their ideology. In both of them, behavior is the terminus of an input-output process. In the cognitive model, the disturbance of water level is predicted from the data, but the controlled variable is not affected until the very end, when the disturbance occurs and the planned action is actually carried out. The arrow from data to assessment actually skips past the four boxes above it so the prediction can be made before the disturbance happens. In both diagrams, the idea that a feedback effect alters the perception during the disturbance, and that the disturbance itself is not perceived at all, is simply missing.

What if the first road had been taken?

Anyone who is convinced of the correctness of either Fig. 2 or Fig. 3 should by now be suffering some doubts. Both of those figures imply a kind of system which, if you tried to build it, would reveal itself to be full of complex calculations and operations. But anyone who decides to accept Fig. 1 even tentatively, just to see what the implications might be, is going to find direct contradictions of important ideas accepted by a very large number of life scientists. That sort of contradiction means either that there is something very wrong with the new idea, or that a revolution has started.

This brings the discussion into the purview of this meeting. How a therapist visualizes what is happening inside a patient or client makes a lot of difference. If Fig. 2 is imagined to be correct, the question becomes that of how to arrange the environment of the patient so as to cause more satisfactory behaviors to take place. If Fig. 3 is imagined, as is likely where a cognitive therapy is envisioned, then correcting a problem becomes one of changing

assessments and predictions of the experienced world and formulating realistic plans of behavior to reach properly defined goals.

But what if Fig. 1 is accepted? What seemed to be environmental stimuli or data for analysis are now just disturbances applied to other variables that are the ones actually under control. The goals are still there for cognitive scientists to find, but now we see that they are goals for perceptions, not for actions, and that the actions are produced and varied in whatever way is made necessary by the disturbances, without any need for complex computations. Behavior is, for the behaving system, relatively uninteresting and unimportant. A person is really concerned about the perceptual consequences of behaving. The behavior that controls those consequences is itself of interest mainly when it affects other people. Clearly a different sort of therapeutic approach is needed if Fig. 1 is the right one.

Conclusions

The water level control system is not complex or hard to understand. The greatest difficulties in assimilating PCT come not from its complexity but from the conflicts between PCT and other theories learned and accepted long ago. That's the main message I want to convey here. There is no way simply to add PCT to the older theories: a choice is necessary. In both of the older views, what an organism does begins in the environment and ends with actions on the environment. Under PCT, the only reason for action is to affect a controlled input to make a perception match an internal specification, a goal state. Seeming stimuli are, in most cases, only disturbances affecting the real stimulus.

There is a conflict now in the worlds of all the life sciences. It is a conflict between the new ideas embodied in perceptual control theory, which are simply the principles that Ktesibios unknowingly put into practice 2200-odd years ago, and the old concepts of what behavior is and how it works, which were developed in the 17th through 20th Centuries because theoreticians failed to rediscover what Ktesibios saw so long ago. Resolving this conflict is probably going to be a long process. Every person now pursuing PCT has felt the inner conflicts, and the resolution is far from finished. One does not dump a lifetime of learning overnight even willingly, and willingness is not an easy state to reach. I hope a few who read and hear this will find it more possible to become willing.

About the Author

William T. Powers is the creator of Perceptual Control Theory (PCT) and the author of several books and papers on that subject. He may be contacted by email at powers_w@frontier.net.