

Date: Sun Dec 11, 1994 6:19 am PST
Subject: Center proposal

[From Dag Forssell (941210 1200)]

>[Bill Powers (941207.0730 MST)]
>To: Fred Good
>CC: CSG-L

Bill's post to Fred Good refers to an outline Bill posted to 16 netters in mid-August. The outline was addressed to Fred Good, publisher of New View Publications, who now distributes what was CSG books. Fred attended the Durango conference in July and has undertaken to locate and approach foundations, soliciting support for a proposed center. Because of the need to update the PCT Introduction and Resource Guide, I have spoken with Fred recently and learned that the research phase of his effort is finished. Fred told me that he identified some 20 foundations and that these things take on the order of two years to come through. Fred is joining CSGnet.

With this reply to Bill's post I quote the letter to Fred in its entirety [but in pieces] and append the proposal with my comments inserted in brackets. I believe the proposal may be of interest to all netters. Parts of it may serve in a revised introduction to CSGnet??? (any volunteer to give it a fresh look?) I hope to see further comments from lurkers as well as regulars on this topic.

> Hi, Fred --

> Sorry about long delay -- I kept putting off writing because I wasn't sure what to say about the proposal for a Center for the Study of Living Control Systems.

> I was wondering what had happened to the proposal, and not only on your side. Since I sent those contributions for the proposal there's been hardly a peep from anyone -- no suggestions, criticisms, or enquiries, despite some direct contacts from me asking for commitments.

Bill, you don't get criticisms on well written material. No disturbance -- no response, remember? With a distribution of 16, no-one may feel particularly called on to respond. Besides, it was understood that this takes time. Also, I suspect you exaggerate. You may have forgotten comments by phone etc.

> The only conclusion I can draw is that there isn't much enthusiasm for starting up a Center. This has made me look at my own real feelings about it, and I realize that I have been rather dreading the thought of going back to work, which is what it would amount to for me. I ran two careers in parallel for 40 years, but I just don't have the gumption I once had. Starting a Center is a young man's job.

Two subjects here. I see no basis at all for your conclusion in the absence of feedback. Your own reluctance to work at it 80 hours a week is another matter entirely.

Let me get philosophical for a moment. I find that as I have internalized the systems concept of PCT, I have no choice but to see myself as an autonomous living control system. I value my own ability to control my perceptions freely. In my roles as father, husband, friend, teacher or manager, I now strive to support others, especially those close to me, to control theirs in a way that is satisfying to them. (This motivates me to teach PCT to the world). PCT has become very important in my life. Once I have learned PCT and decided that it makes more sense than any other explanation for life that I have ever considered, I cannot erase it. I expect PCT to be part of me for the rest of my life.

I think it is very important to teach PCT to the world. The impact on society can be far-reaching. For instance, I see many practices change in management of organizations as people understand PCT and change their personal values as a consequence of that understanding.

The idea of a Center for the Study of Living Control Systems is a systems concept that fits with the understanding and values I have as a consequence of learning PCT. The idea can be implemented in many different ways, on many different scales. The idea does not depend on any one individual, not even Bill Powers. I think that if some foundation gets the idea and wants to fund it two or three years from now, some individuals who understand PCT will step up to the plate, as committed to PCT and its importance as PCT enthusiasts you can name are today. But I do not think it necessary to name the team today. People who understand PCT will be there. In the meantime, the effort to solicit support from foundations -- people with high ideals, connections and resources - may pay unexpected dividends. We talk of how people are self-selected and come to PCT. As our recent exchange on the net (ending with my post on selling PCT right) shows, we seldom have discussed (from the first principles of PCT--not miscellaneous personal opinions) how to make it possible for people to select themselves. I see this effort as yet another way to make it possible for people (with some clout) to select themselves.

As I write this, the thought comes to me that I have no doubt that the world will eventually benefit from PCT. The idea of the center is to accelerate that. To begin to solve the pressing problems of society ten years from now, not forty. If this point is valid, it should be made very clear in the proposal, since funding would make that difference.

- > The idea of the Center is great, but it's probably premature. PCT has not reached a degree of acceptance such that the scientific and academic community would see the Center as legitimate. Whatever we may think about the attitudes of today's conventional scientists, colleges and universities are still the intellectual centers everywhere in the world. Without that whole community there's no infrastructure, no large-scale access to students, no contacts with the power structure, no influence on intellectual or educational leaders, no continuity with history.
- > I don't want to see the PCT movement develop at the level of Glasser's Institute for Reality Therapy. However big that empire, it's still basically an amateur operation cut off from the rest of the world except through fringe contacts. No, I don't mean "amateur," what's the word I'm looking for? It's less like a university and more like the Nairopa Institute or the Ayn Rand Institute -- a private enterprise not subject to the discipline and challenges of the academic world in general, and far too limited by the whims, prejudices, and beliefs of a few people.
- > God knows I've become fed up with academia and establishment science, but that's the real world in which human knowledge is advanced, however slowly and imperfectly. It is, truly, the mainstream of human thought even though it contains a lot of backwaters and eddies going the wrong way. We have to look at the other side of the story: what the mainstreams of science have accomplished. When there is a consensus on a good idea, amazing things happen. If we want to contribute to the future of the human race, we simply have to keep working toward a consensus. Splitting away and starting our own little universe will have just the opposite effect.

The idea of the center is one thing. The form and scope of its organization is another. You make several good points here. But nothing precludes a foundation funding a center associated with or part of a legitimate university. -- Just one possibility. This is something that can be determined in discussion with the foundation and its contacts two years from now.

- > I'm especially leery about starting up a separate universe when the very concept of science is being rejected by people in such large numbers. There seems to be an almost deliberate return to ignorance going on in this country. Enormous numbers of people seem to have lost their skepticism. There seems to be a movement back toward superstition, magic, and fairy tales, a yearning to return to an imagined past that never did exist.

I agree with you. I find the superstition in our society scary.

> I don't know; I sound to myself like a crotchety old man. But right now I think that the time to start the Center is when events in the world of science demand it. Despite what I said in the writings for the proposal, we have not quite yet been overwhelmed by people clamoring to know more. Maybe we just have to wait for that to happen; maybe then there will be people in the PCT movement who can see enough of a future in it to risk making a commitment. In the meantime, we're not doing too badly.

Events in the world of science will never "demand" PCT. We will never be overwhelmed with people clamoring to know more UNLESS we make it possible for people to imagine how their lives can be better from the study of PCT. If that means suggesting that you can figure out what it takes to resolve conflict with mutual satisfaction from a study of PCT, then that may be what we should say. I believe PCT shows the way to much more than that, though. See section III in the proposal below.

I believe that rather than drop the project, we should charge ahead with appreciation for the work Fred Good has done so far, and support his efforts to contact foundations with an edited proposal -- a rearrangement, refinement or replacement of the draft below.

CSGnet lurkers: What do you say? Who has experience relevant to this proposal? What does PCT itself tell us?

Date: Aug 12 - 19, 1994
Subject: Center proposal. Draft by Bill Powers.

Assembled with comments by Dag Forssell. December 10, 1994.

- I. Introduction
- II. The mission of the Center
- III. Assessment of needs: why the life sciences need PCT
- IV. Objectives: How needs will be met
- V. Methods: What we will do to accomplish the objectives
- VI. Evaluation: How we will make sure we're doing what we propose to do.
- VII. Future funding: plans beyond the grant period
- VIII. Budget: what funds we need, what other sources will be sought.

Proposal for the organization of a
Center for the Study of Living Control Systems

I. Introduction

In late July of 1994, the tenth annual meeting of the Control Systems Group took place at Fort Lewis College in Durango, Colorado. For ten years, members of this interdisciplinary group have been exploring, improving and teaching a new general theory of behavioral organization called Perceptual Control Theory. The fields of application have been diverse; they have included biochemistry, clinical psychology, control engineering, economics, educational theory, experimental psychology, linguistics, management consulting, marriage counselling, molecular biology, neuroscience, organizational design, philosophy, physiology, psychotherapy, rehabilitation, school reorganization, and sociology. Scientists and other professionals in all these fields have found in Perceptual Control Theory not only a common language that communicates across the boundaries of their specialties, but a fundamentally new insight into the nature of behavior that seems to throw new light on essentially every branch of the behavioral sciences.

The work leading to PCT began in the early 1950s at the V.A. Research Hospital in Chicago, Illinois. Two physicists, William T. Powers and Robert K. Clark, and one clinical psychologist, the late Robert L. MacFarland, set out to explore the concepts of negative feedback control as put forth by Norbert Wiener, Arturo Rosenbleuth, and Julian Bigelow in their pioneering book, Cybernetics: control and communication in the animal and the machine. This happened during a brief window in history when the analog computer was the most widely-used tool for simulating living systems. Soon the digital computer

took over, leading the mainstreams of behavioral science (and even Cybernetics) into a very different approach to the function of the brain and the nature of behavior. But the original three researchers persisted, publishing their first work on a "General feedback theory of human behavior" in 1960 -- just in time to see the scientific world swinging onto a new path and declaring control theory out of date.

The original group disbanded, but Bill Powers continued to develop the basic ideas, lecturing occasionally and publishing a few papers, encouraged by a few people such as Donald T. Campbell who urged him to continue the work and keep it alive. In 1973, Powers published Behavior: the control of perception, which attracted a much wider audience. Here and there, scientists began to pick up on the basic concepts, often astonished to find how long they had been in existence and how difficult it had been to achieve acceptance of them among conventional behavioral scientists. Year by year as Powers continued to publish, lecture, and give seminars at psychology departments, the number of serious adherents grew, until in 1985 there were enough interested people in one place at one time to decide to form the Control Systems Group, or CSG.

In the ten years of its existence, this group has grown from 15 members to 60 scattered across the United States and Canada. About four years ago, a discussion group was begun on the internet; the list now includes around 130 residents of 19 countries, only perhaps 20 of whom are part of the original CSG. As a result of this wider discussion, the first European Workshop on Perceptual Control Theory was held in June, 1994 under the auspices of the University of Wales in Aberystwyth, attracting participants from England, Wales, Scotland, France, and Germany, as well as members of the American CSG. The next meeting, as well as research proposals for cooperative work in prosthetics and rehabilitation medicine in the UK and United States, are now being planned. As this is being written, paper collection is underway for a special issue of the International Journal of Human-Computer Systems devoted entirely to PCT and edited by the senior psychologist of the Defense and Civil Institute for Environmental Medicine of Toronto, Canada. In the United States, a special two-day seminar on PCT is being planned (by a Dean of Education at an Eastern University) for the next meeting of the National Educational Research Association. A prominent and widely-cited cellular biologist is organizing a new institute for biochemical research and has announced that he plans to use PCT as a major organizing principle for the work of that institute. In Houston, Texas, a program for using PCT to assess performance before and after surgery for spinal injuries is under development; the neurosurgeon in charge has postponed all major surgeries until the new system is in place. In Phoenix, Arizona, a group of school principals and other administrators has been applying PCT to discipline problems in inner-city schools, with remarkable results. A member of the CSG, Edward Ford, has been travelling around the country by invitation, teaching teachers, parents, and school boards how to apply the principles of PCT. There seems suddenly to be a ferment of interest in PCT spreading more rapidly and more widely than ever before in its 40-year history.

All this has come about without any plan or any organized attempt to gain publicity -- all but accidentally. One person has taught another. All the members of the CSG and of the internet group selected themselves, having heard or read something about PCT and wanting to know more. There are signs, however, that this unorganized approach will not suffice much longer, and that the demand for education and research in PCT will very soon exceed the informal capacities now available. That is why we are planning a Center for the Study of Living Control Systems, and why we are asking support to make it a reality.

["will not suffice much longer" -- suffice for what? This paragraph needs work. Concept: will not suffice to teach enough people to make an impact, but will be held in check by institutional inertia, delaying the spread of ----. It may suffice for a 40 yr adoption, not for a 10 yr adoption.]

II. The mission of the Center

The Center for the Study of Living Control Systems will serve to organize the teaching of PCT, to coordinate research on PCT at colleges and universities in the United States and abroad, to conduct scientific meetings, workshops, and

practicums on PCT, and to plan curricula leading to advanced degrees in which PCT plays a major part. It is hoped that the Center will be able to provide facilities for visiting scholars and researchers where members of diverse disciplines can come together to do real work on the nature of living control systems. It is hoped, too, that the Center can provide public education on the meaning of PCT in practical affairs, in industry, in the management of organizations.

By the nature of its subject matter, the Center will have, we hope, a limited lifetime. Perceptual Control Theory is a theory of life; it is a tool to be used wherever the phenomena of life are studied. The primary purpose of the Center is to sharpen this tool to make it as useful as possible, to put the tool into the hands of knowledgeable users, and to pave the way for mainstream research to take over the use of this tool. At some point in the future, in 10, 15, or 20 years, PCT will be taken for granted in all the sciences of life. There will then be no more need for a Center for the Study of Living Control Systems than there is for a Center for the Study of the Roundness of the Earth.

III: Assessment of needs.

The fields of psychology and social science, which we can call collectively the behavioral sciences, are, in the opinions of many of their own practitioners as well as others, in deepening trouble. One indication of the basic problem can be seen in a published survey of the journals: the median correlation found between manipulated variables and the behaviors they supposedly explain or predict was only 0.26.

[How about a reference for netters, perhaps for foundation. June 29, 1995: I found it. Runkel mentions it in "Casting Nets" on page 82-83. It is Daniel J. Brown: Mirror, Mirror... Down with the Lineal Model. American Educational Research Journal. Fall 1975, Vol. 12, No. 4, Pp491-505.]

Just how bad this figure is can be conveyed by a seldom-used statistical measure called the "coefficient of alienation," known informally as the coefficient of uselessness. If r is the correlation, the coefficient of uselessness is the square root of $1 - r^2$. For a correlation of 0.26, the coefficient of uselessness is 0.966. In approximate terms, this means that if a supposed causal variable is manipulated and is discovered to have a correlation of 0.26 with some measure of behavior in a population of subjects, a prediction of the behavior of any one subject based simply on the average behavior of the population would be about 97 percent as accurate as a prediction based on knowledge of the manipulated variable. So the manipulated variable is about 97 percent useless as a way of predicting the behavior of individuals, regardless of how "significant" the correlation is. Half of the studies necessary came out with lower correlations than 0.26, which was the median.

[Math to involved for a foundation executive??]

The basic problem is that the theories of the psychological and social sciences, the behavioral sciences, don't work for individual human beings. Where they do work to some extent is for organizations like governments, insurance companies, advertising firms, hospitals, and schools where the measure of success is based on population averages and the welfare of any one individual is irrelevant to the track record. Because of this emphasis on mass statistics, it is possible to show small but provable benefits of statistical measures to the organization, yet at the same time to show that it is to no individual's advantage to be judged by the same measures. In terms of the average correlations found in the literature, the individual has nearly an even chance of being described and treated incorrectly, often at great cost to the individual.

The failure of psychological theory has led to a split within psychology itself. The "scientific" psychologists and the clinical psychologists no longer even talk to one another. They have separate journals, separate meetings, and separate societies that do not interact except for mutual recriminations. Theory and practice in psychology have essentially nothing to

do with each other. And psychology has no relation to sociology, and so forth through the whole set of independent and unconnected subdisciplines. The whole enterprise is coming apart.

One of the great errors in the behavioral sciences has been to assume that the importance of a problem somehow makes up for the inadequacy of the methods being used to try to solve it. Certainly war, poverty, ignorance, mental illness, the breakup of families, and crime are serious problems and we need to solve them. When the impossibility of solving them using knowledge based on correlations of 0.26 is pointed out, the rejoinder is that it's better to do something with what we have than to do nothing. But the claim that using such methods is "doing something" is warranted only in the sense that there is some kind of activity going on rather than none. And even worse, because we are "doing something," the implication is that we should continue doing it some more -- not because it is doing any good, but simply because we are doing it.

What is needed is a basis for behavioral science that is as solid as the basis of physics. The only problem with establishing such a foundation is that we would actually have to go back to basics, to do experiments as simple as those of Galileo rolling little balls down an inclined plane and timing their speed with his pulse. Galileo had no inkling of the physical sciences that would grow out of his simple determinations of properties of nature. He was not trying to solve the big problems, but the little ones. The reason he had to start with the little problems is that he was looking for answers that would remain true under all conditions, all of the time; not correlations of 0.26. Only when a collection of predictions of very high quality existed could he or anyone else go on to construct the next level of complexity, and the next. Only by maintaining the standard that a theory should explain observations as accurately as we can measure them could this progression be maintained from year to year and century to century. Only in that way could a science come into existence: out of small facts of very, very high quality.

The most ambitious undertaking that could be mounted in the behavioral sciences would be to start all over from scratch. This would mean questioning all the old assumptions, looking for ways of characterizing behavior which, no matter how simple or seemingly unimportant, would correctly predict behavior to the limit of our ability to measure it. That how physics and chemistry got started. It is the only way in which a true science of behavior can be built.

IV: Objectives: meeting the needs

In the Introduction, the history of the Control Systems Group, and of the theoretical approach behind it, was laid out. From that, it should be clear that work on a basically new science of behavior was actually begun long ago. As in the case of many scientific revolutions, it was made possible by a discovery in an unrelated field: the discovery by electronics engineers of the principles of negative feedback control. It is seldom remembered that this discovery came about because the engineers were trying to build devices that would imitate a living system: a person performing the kind of behavior we call controlling.

As in many other cases, the real significance of the discovery was not immediately recognized. The founders of PCT were not aware of the radically new nature of the control-system model; indeed, they naively thought that once it was developed to a certain point, behavioral scientists would welcome it as a step forward. Twenty years passed before the realization came that control theory and the phenomena to which it is addressed are completely at odds with the fundamental assumptions behind the behavioral sciences. The inexplicable refusal to treat these new ideas seriously suddenly became understandable: if the control-system model was right, then essentially everything important that the behavioral sciences thought they knew about behavior was wrong. To adopt control theory was for all practical purposes to start the behavioral sciences over again, from scratch.

The small number of people actually doing experimental work is a partial explanation of why progress has been so slow, but is by no means the full explanation. What has really made progress so slow is that experimental work has been done under a very severe requirement: that every basic concept,

before being accept as part of the "official" model, be demonstrated beyond doubt, quantitatively. Finding the right instance of a PCT model was a matter of rejecting many models which did not meet this criterion.

It is very easy to do statistical research; practically any hypothesis, if tested on a large enough population, will eventually produce correlations with a publishable confidence level. But if the correlations one insists upon getting out of an analysis (of any kind) are 0.95 and up, the number of successful experiments must inevitably drop very sharply. It is not easy to find phenomena of nature that are highly reliable and replicable. It is no easier in psychology than it is in physics. But it is just as necessary in either science.

Real experimentation with PCT was not possible until small computers became available to workers operating on small or nonexistent budgets. The first tentative experiments with PCT were done with a big slow minicomputer programmed with paper tape, and then with a much more powerful home-built computer kit with a memory capacity of 0.016 megabyte. Prior to that time, only possibilities could be explored using pencil-and-paper analysis: there was never any funding for experimental equipment.

What seems obvious to us now was, in 1973, a complete mystery. It was not self-evident how a model of behavior should be constructed, even though a hypothetical outline of the structure existed. Once the small-computer revolution made the design of real-time experiments feasible, many very complex possibilities were explored. There was a reluctance to believe that the correct model could be simple. Yet the best model of simple tracking behavior that was first found could be written in four short lines of Pascal computer code:

```
e := r - c
c = d[t] + h
h := h + k*e
t := t + 1
```

[This part is probably too technical.]

This code segment, executed again and again, is a model of tracking behavior in a human being. The human being uses a handle to hold a cursor stationary at a fixed target location; the cursor is being pushed this way and that by a continuously varying disturbance (d[t]). For any random pattern of disturbance that is smoothed to a bandwidth of about 1 Hz, which makes holding the cursor in one place fairly difficult, this model will create the same pattern of handle movements that the live subject creates for the same disturbance pattern, with an accuracy of about 5 percent. If a tracing of the subject's handle movements against time is laid over a similar tracing created by the above program, the two traces will coincide everywhere, for a 60-second run at 30 samples per second, with a mean square error of five percent of the peak-to-peak variations, or less. The correlation between the model's simulated handle behavior and the subject's real handle behavior is typically 0.99 or higher.

This degree of match of model to real behavior is achieved by adjusting a single constant: k in the third line above. For a fully practiced subject, k comes to a very stable value, which repeats experiment after experiment.

Using this best-fit value of k, it is possible to run the model with a new pattern of disturbances to create a prediction of a subject's handle movements in a later experiment employing the same new disturbance pattern. When that is done, the predicted behavior still correlates above 0.99 with the subject's behavior in the subsequent test. And, as a CSG member with extraordinary patience proved, a prediction made for a particular subject in this way, using a pattern of disturbance not before experienced, achieved the same level of accuracy when the experiment to test the prediction was done five years later. The five-year test was done at a CSG meeting. When the results appeared on the screen, there were cheers. In another five years, the 10-year test will be done, using a different disturbance pattern and prediction generated at the same time as the first one. There are three more data sets waiting for the remaining three tests.

The behavior that is the subject of these experiments is considered simple even by PCT researchers; they call it, affectionately but realistically, "stick-wiggling." It is a simple and unimportant kind of behavior. But that is not what matters. What matters is that we have a model and a number characterizing real individual human behavior which together constitute a fact of the very highest quality, at least an order of magnitude better than any fact that the behavioral sciences have been able to produce before. This first fact is a foundation stone, and when enough such foundation stones have been laid, a science can be built on them.

Since the first stick-wiggling experiments, many more have been done. Human control of shape, size, relationship, rotation rate, and sound pitch have been done, with the same or a slightly more advanced model (containing perceptual delays) being matched to them with the same accuracy as initially found. More high-quality facts, more stones in the foundation. There have been tentative explorations of even more complex behavior, for an extreme example the behavior of defending a self-concept against invalidation. Here no quantitative measures of handle positions could be used, but when the right experiment was found (after several false starts), the number of subjects demonstrating exactly the predicted phenomenon was 25 out of 26; one response was unscorable.

In the realm of modeling, a PCT simulation of crowd behavior was done, in which each individual was a set of two or three elementary control systems controlling for one or more specific conditions: reaching a target, following another person at a specific distance, and avoiding collisions with other people or stationary objects. Up to 256 active persons could be on the computer screen at one time. Various setups were tested, producing phenomena which a sociologist immediately recognized as recurrent phenomena seen during his decades of observational work on crowd behavior.

Still another model translates the stretch and tendon reflexes in a human arm into a PCT type of model, which explains clearly how these reflexes work. Combined with a model of visual control systems, this "Little Man", which lives in a desktop computer, can reach out and track a moving target in three dimensions, in real time, showing trajectories like those seen in real human subjects. The response of the arm to muscle forces is calculated using a correct model of physical dynamics; a realistic model of muscle operation is used. Models offered by others to explain the same behavior are too complex to work in real time; they involve the calculation of inverse dynamics and kinematics which are not required in the PCT model, and they cannot work in the presence of unpredictable disturbances as the Little Man can. And the Little Man can not only reach out and touch the target; he can draw continuous circles in the air around it, even as it moves in arbitrary and completely unpredicted patterns.

The crowd simulation and the Little Man program are the first signs of a skeletal structure beginning to rise from the foundations.

Everything that the Center for the Study of Living Control Systems does will be aimed at extending the foundational work and exploring the first layers of the experimental science that will rest on the foundation. A great deal of additional work will be needed and many people will be needed to carry it out. So we must raise a new generation of scientists prepared to do this research, and attract as many established scientists as possible into it, wherever they choose to do their work. We must teach this new collection of researchers to maintain a standard for the acceptance of new facts that is very different from the one that is presently maintained in the behavioral sciences. The fact that progress can be made under the most exacting standards, however slowly, should encourage others to try doing the same thing. The results will stand for generations.

Comment by Bill P. with transmission of section V.:

Here is a possible next section in Fred Good's outline for our proposal. It introduces some new considerations that may suggest rewrites of the preceding sections. Please contribute to this and say what you think of it. If Fred pulls this off, the pipe dream may suddenly start looking like

a reality, and call for some concrete decisions and commitments from those who claim to be supporters of PCT. We all have to start thinking of what we would really do if suddenly \$2,000,000 appears with the go-ahead to carry out the plan. Life could not continue as it is now.

V. Methods.

The transition from the current state of research in PCT to a fully functioning Center will require achieving a series of intermediate goals as well as a change of focus. Those interested in this field have, over several decades, become used to working in relative isolation most of the time, on small or nonexistent budgets, in whatever time could be spared from the requirements of making a living and a career. Those who have worked the longest are nearing or have passed retirement age, and while still productive are no longer driven by the ambitions and hopes of youth, or equipped with its energy. When the Center reaches the point of becoming fully operational, it should be under the command of a younger generation with the elders serving in an advisory capacity.

The first business of the Center, therefore, must be that of passing the torch so that all that is of value in the knowledge and lore of PCT might be preserved in the generation that will carry on the work into the 21st Century.

The first step: internal education

As matters stand now, there are perhaps six or eight people who understand both the general principles of PCT and the detailed mathematical and practical approaches to systematic quantitative research in this field. Following closely behind them is a larger group which has a good grasp of principles, but little experience with modeling and experimentation of the kind needed. And the largest group of those interested has only a subjective and practical knowledge of PCT, with a limited ability to apply it in new areas. This situation is the result of the realities of life: a group that meets only once a year has little chance for systematic teaching and learning, so that progress up the ladder of understanding is very slow. There is no basic textbook, there is no curriculum of background studies, there is little time for concentrated learning, and for younger members there has never been any way to devote full time to the study of PCT while developing a career and supporting a family.

The first step toward creation of a Center for the Study of Living Control Systems must therefore be to bring together a substantial number of those interested and turn them into experts with the skills and knowledge required to initiate and carry out PCT research -- or at least put them well on the road toward this end. A core group of ten to twenty people is needed. These people will be required to go through considerable inconvenience and to make a commitment that may have repercussions on their careers and income, because the training period should go on for the equivalent of at least a full school year.

Our of this initial training period there should emerge a new generation of PCT researchers ready to teach others and to take up the threads of research that have been started during the last 20 years. The temporary college may well become a permanent one, the physical and intellectual nucleus of the Center. The experiences of this year of teaching will lead to the writing of a textbook and the creation of a curriculum which can be adopted by other centers of education. The infrastructure developed during this year will become the mechanism for coordinating research and applications, for organizing interdisciplinary meetings both national and international, and for organizing applications of PCT in schools, businesses, and other places. The physical plant will become a place for visiting scholars and researchers to carry out projects and establish collaborations. In short, we should emerge from this initial year with a Center that is nearly in operational condition.

It should be mentioned that the "students" who will take part in this project will not be any ordinary kind of students. Most of them will have degrees in some discipline, and some will have had ten or twenty years' experience as university professors, or the equivalent in some profession. We will encourage the participation of some new college graduates or candidates for advanced

degrees, but the initial cadre will be largely composed of experienced people -- who happen to have lacked the particular training needed for a technical understanding of PCT. The lack of this training can be understood if it is realized that these people will come from fields such as educational research, linguistics, sociology, biochemistry, economics, clinical and experimental psychology, and organizational development. Although the emphasis will be on filling in the technical aspects of an education in PCT, the course will greatly benefit from the experience of the students, so in some respects it will resemble a workshop. These rather overqualified students will be well-equipped to write a textbook on the very subjects they have just finished assimilating, and to make it an interdisciplinary textbook as well.

The sponsor is going to have to be firmly convinced of the long-term value of PCT research, because a considerable financial outlay will be needed to make this project feasible. In effect, we will be setting up a temporary college, which will require a physical location, equipment, and administrative support. We may need to hire teachers -- for example, a mathematics teacher, a teacher of basic servomechanism technology, and a teacher of computer programming. It will be necessary to support some or all of the attendees as well as those teachers without an independent income with a sufficient salary to maintain their families while this educational period is going on. And preparing the curriculum is going to take on the order of six months, during which time the teaching group will require support and a place to work.

Furthermore, the real test of the Center will come only after the first year; the initial group of students will turn around and become those in charge of the Center and the teachers of the next group, which will surely include many applicants from outside the Control Systems Group. As students return to their institutions, there will be increasing interest in learning about PCT; it will quickly become important to provide curriculum materials to colleges and universities if only to avoid swamping the Center.

So the total support that the sponsor must contemplate providing will have to cover not just the first year, but two or three years following, perhaps for a total of as much as five years. After the first year or two, the Center may be able to generate income through charging for the main courses, and for consulting with schools, businesses, and governments concerning problems to which PCT may offer solutions. But the sponsor should consider that this must be a relatively long-term commitment.

[No draft has been prepared for the following three sections. As I understand it, the breakdown in eight sections was suggested to Bill by Fred Good, as appropriate for a generic proposal. There is probably nothing sacred about this outline. Fred tells me that foundations have widely divergent expectations of proposals. I think the idea of PCT and the idea of teaching it sooner rather than later is the central part of our proposal. The scope and form of a center can be highly flexible.]

VI. Evaluation: How we will make sure we're doing what we propose to do.

[Is this important as part of the proposal? Whoever funds the center can participate and be fully informed. The idea of PCT, once conveyed, will make the funding decision maker committed.]

VII. Future funding: plans beyond the grant period

[Is this important as part of the proposal? Do we have to answer up front? Again, the idea of PCT, once understood, might make it moot. This kind of issue, if it is an issue, might be better discussed with a real, flesh and bones, funding person who has particular concerns or constraints.]

VIII Budget: what funds we need, what other sources will be sought.

["what funds we need" is a problem. We don't. A funding person who understands PCT might be prepared to commit \$250,000 per year. We can do something with that. A funding person who can see the potential of PCT to promote world peace might want to commit \$10 million per year for 10 years. We can do something with that. All kinds of life scientists will suddenly participate and publish up a storm. Money talks.

I find it hard to get excited about these last sections. I imagine that is why Bill has not offered a draft in response to Fred's request. How about leaving these open. Let us concentrate on developing a proposal that is compelling in the first sections. Perhaps the last sections can be glossed over to almost nothing.]

Best, Dag