ADAPTIVE CONTROL OF ILL-DEFINED SYSTEMS

NATO CONFERENCE SERIES

I Ecology
II Systems Science
III Human Factors
IV Marine Sciences
V Air–Sea Interactions
VI Materials Science

II SYSTEMS SCIENCE

Recent volumes in this series

Volume 5 Applied General Systems Research: Recent Developments and Trends
Edited by George J. Klir

Volume 6 Evaluating New Telecommunications Services
Edited by Martin C. J. Elton, William A. Lucas, and David W. Conrath

Volume 7 Manpower Planning and Organization Design
Edited by Donald T. Bryant and Richard J. Niehaus

Volume 8 Search Theory and Applications
Edited by K. Brian Haley and Lawrence D. Stone

Volume 9 Energy Policy Planning
Edited by B. A. Bayraktar, E. A. Cherniavsky, M. A. Laughton, and L. E. Ruff

Volume 10 Applied Operations Research in Fishing
Edited by K. Brian Haley

Volume 11 Work, Organizations, and Techonological Change
Edited by Gerhard Mensch and Richard J. Niehaus

Volume 12 Systems Analysis in Urban Policy-Making and Planning
Edited by Michael Batty and Bruce Hutchinson

Volume 13 Homotopy Methods and Global Convergence
Edited by B. Curtis Eaves, Floyd J. Gould, Heinz-Otto Peitgen, and Michael J. Todd

Volume 14 Efficiency of Manufacturing Systems
Edited by B. Wilson, C. C. Berg, and D. French

Volume 15 Reorienting Health Services: Application of a Systems Approach
Edited by Charles O. Pannenborg, Albert van der Werff, Gary B. Hirsch, and Keith Barnard

Volume 16 Adaptive Control of Ill-Defined Systems
Edited by Oliver G. Selfridge, Edwina L. Rissland, and Michael A. Arbib

ADAPTIVE CONTROL OF ILL-DEFINED SYSTEMS

Edited by

Oliver G. Selfridge

Edwina L. Rissland

and

Michael A. Arbib

University of Massachusetts
Amherst, Massachusetts

Published in cooperation with NATO Scientific Affairs Division

PLENUM PRESS · NEW YORK AND LONDON

Library of Congress Cataloging in Publication Data

NATO Advanced Research Institute on Adaptive Control of Ill-defined Systems (1981: Moretonhampstead, Devon) Adaptive control of ill-defined systems.

(NATO conference series. II, Systems science; v. 16)
"Proceedings of the NATO Advanced Research Institute on Adaptive Control of Ill-defined Systems, held June 21–26, 1981 in Moretonhampstead, Devon, England"—T.p. verso.
Bibliography: p.
Includes index.
1. System theory—Congresses. 2. Adaptive control systems—Congresses. I. Selfridge, Oliver G. II. Rissland, Edwina L. III. Arbib, Michael A. IV. Title. V. Title: Ill-defined systems. VI. Series.
Q295.N37 1981 003 83-17699
ISBN 978-1-4684-8943-9 ISBN 978-1-4684-8941-5 (eBook)
DOI 10.1007/978-1-4684-8941-5

Proceedings of the NATO Advanced Research Institute on Adaptive Control of Ill Defined Systems, held June 21–26, 1981, in Moretonhampstead, Devon, England

Softcover reprint of the hardcover 1st edition 1984
A Division of Plenum Publishing Corporation
233 Spring Street, New York, N.Y. 10013

PREFACE

There are some types of complex systems that are built like clockwork, with well-defined parts that interact in well-defined ways, so that the action of the whole can be precisely analyzed and anticipated with accuracy and precision. Some systems are not themselves so well-defined, but they can be modeled in ways that are -- like trained pilots in well-built planes, or electrolyte balance in healthy humans.

But there are many systems for which that is not true; and among them are many whose understanding and control we would value. For example, the model for the trained pilot above fails exactly where the pilot is being most human; that is, where he is exercising the highest levels of judgment, or where he is learning and adapting to new conditions. Again, sometimes the kinds of complexity do not lead to easily analyzable models at all; here we might include most economic systems, in all forms of societies.

There are several factors that seem to contribute to systems being hard to model, understand, or control. The human participants may act in ways that are so variable or so rich or so interactive that the only adequate model of the system would be the entire system itself, so to speak. This is probably the case in true long term systems involving people learning and growing up in a changing society.

Some kinds of true complexity arise when simple systems are put together into numerous and complex assemblies; one might hope that thermodynamics showed the way. But, alas!, the simple truths that thermodynamcis seizes from simply behaving molecules may well be unique in science; most other examples -- from schooling of fishes to economic models of simple consumers -- are just too intractable when they are built on any basis that corresponds roughly to reality.

One kind of difficulty seems to arise inevitably from considering purposive behavior, whether directly by people or imputedly by complex mechanisms: that is the interaction of different purposes and goals. Sometimes the goals are in immediate conflict, so to speak, sometimes they are cooperative, but in real systems they are always embedded in hierarchical structures of goals subsumed under the values of the designers and operators.

Nevertheless, although we cannot attain with some systems a degree of understanding and control that we might with others, that does not mean that we cannot strive to find out what controls are possible. The notion of control here does not necessarily mean just the dimensions of decision and manipulation; rather it also includes the concepts of improving understanding, gathering data, and making inferences about the feasible structures of affectors and effectors. That is, we are interested also in the science, as well as in the technology.

The techniques of adaptive control are well enough in hand for the large class of dynamic systems where the plant and its dynamics can be analyzed <u>a priori</u> using the mathematical techniques of modern control theory. The purpose of this volume is to evaluate the control techniques that may be applicable when the system is not well defined, or may even be frankly unknown; such systems would include, <u>inter alia</u>:

human control systems where the person is learning
complex communication networks, such as those involving people
biological systems
economic systems
man-machine systems

No general rules are offered that will enable control in all, or even most, such systems. Rather, we analyze the state of the art, propose new approaches, and discuss possible applications and implications.

Most of the papers were presented at a NATO Advanced Research Institute held in Moretonhampstead, Devon, England, from June 21 to June 26 of 1981 to analyze the current technology and the current problems, and to assess the possiblities and to formulate a program. For the purpose of this volume, the papers have been grouped into six sections. Section 1 offers a variety of "Perspectives": a dialogue constructed by the Editors to dramatize a number of the issues addressed throughout the book; perspectives from the point of view of man-machine system, adaptive algorithms, and economics; as well as a historical view of adaptation in biology. Section 2 looks at adaptation and control using the well-defined mathematical techniques of control theory and nonlinear analysis, in contrast to Section 3 in which the planning techniques of Artificial Intelligence and the power of the computer take over from the closed forms of the mathematician. The final 3 sections look at different application areas in which issues of adaptation are raised with respect to ill-defined systems; "Motor skills" in Section 4, "Language Acquisition and Adaptation" in Section 5, and "Development and Evolution" in Section 6.

We thank Dr. Bayraktar of the NATO committee for making the Advanced Research Institute possible, Devon for being so beautiful in June of 1981, and our authors for all they contributed to the Institute and this volume. We also thank Barbara Nestingen for her invaluable editorial assistance.

Massachusetts, March 1983

Oliver G. Selfridge
Edwina L. Rissland
Michael A. Arbib

CONTENTS

A DIALOGUE ON ILL-DEFINED CONTROL

Michael A. Arbib, Oliver G. Selfridge and
Edwina L. Rissland

Department of Computer and Information Science
University of Massachusetts
Amherst, MA 01003 USA

To give an introduction to the concepts addressed in the present volume, the Editors have contrived a dialogue between two scholars, imaginatively named A and B, which reflects the current state of our inter-discipline. It captures many of the views expressed in general discussions throughout the meeting.

A. When we speak as technicians seeking to understand control strategies, a system is a "them," as distinct from the controllers who are an "us"—dividing the universe into "us" and "them." The task of control theory is to provide an algorithm for how "us" is to control "them," and such an algorithm must be based on criteria of optimality or satisfiability, but the appropriate criteria are often unclear in an ill-defined situation. For example, in writing a technical paper we have the conflicting criteria of "keep it short" and "provide the reader with background material." Very often we do not have a single criterion for our control strategy. The different criteria may even be mutually inconsistent. Of course, if we had a well-defined criterion for trading off amongst these criteria, then we could optimize on that weighted criterion function. But, in general, we do not.

Thus, one aspect of our concern with ill-definedness is not adaptive control of ill-defined systems, but ill-defined criteria for control of systems.

B. My most horrifying example is based on the fantasy that finally we have incredible computers which control the economy and production and the weapons system and everything else, and we finally determine our ultimate criterion and issue the command: "Ensure that there is no more human unhappiness." The response—nuclear war: all human life is destroyed; and there is no more human unhappiness.

But the fact that multiple criteria pose problems doesn't mean that they are ill-defined. Isn't it a paradox to define what you mean by "ill-defined"?

A. That's like saying you can say nothing true about falsity. To speak of a system is to believe that we can limit that portion of the world (spatially or informationally or whatever) that is relevant to our interaction with respect to some stated problem. Given that, we try to come up with enough information about that portion to enable us to make some decisions about how to select or influence our interactions. The "ill-defined" results from the inevitable mismatch in how much we choose to consider of the world around us; and how we choose to describe and represent it. One might almost say that a system is ill-defined if adaptive control doesn't work very well. If adaptive control of an ill-defined system "works" then the system is sufficiently well defined.

B. But you can make a model well-defined by including more and more variables to get closer and closer to the essential interactions of the chunk of the world involved.

A. Hardly. You don't model a gas in terms of atomic interactions. You have to find the right thermodynamic variables. You have to find the appropriate generalizations. In the same way, an economic model with thousands of variables would be better defined if you knew how to pick just a few and monitor their interactions.

B. It's not practical to limit an economic model to just a few variables. There will be many variables that the Treasury and consumer groups will want to monitor even if they do not help explain part of the economy. If you counted all the outputs, you might have many more outputs than you are actively controlling. Also, because there is so much noise in the system, we need to keep track of a host of variables to take advantage of unexpected significances in their correlated trends.

A. Perhaps variables are even the wrong "atoms". Maybe our descriptive building blocks ought to be procedures and little chunks of knowledge.

B. But how can you describe what you can't define?

A. You and I do that all the time but with different vocabularies.

Furthermore, it is not clear that we can define functional interactions between so many variables in a fashion relevant to the adaptations we are searching for. Anyway, I pose as our fundamental paradox that, by definition, adaptive control of ill-defined systems is unsuccessful control of a system—because a system is only defined with respect to a task, and if we succeed in a task, then the system is sufficiently well defined.

B. You seem to be saying that a system is well-defined only if adaptive control works for it, but there seem to be logical systems in which you have a good understanding of what is going on, and yet adaptive control isn't even relevant. First order predicate calculus is well defined, but what is to control there? You don't want to say, I think, that well-defined means control at all by this technique.

A. The last part of your comment suggests a fundamental misunderstanding, namely, that one can take something like the predicate calculus and call it a system. The notion of system that I had thought we were using was the notion of some chunk of the universe changing over time, while the control problem was to influence the trajectory of that system. Therefore, although people may formally use a phrase like "logical system" and say predicate calculus is a logical system, it is not the form of system we are concerned with here.

B. But take my example seriously. Let's assume that we've got this definition of the predicate calculus, but our criterion, following your earlier comment, is ill-defined--namely, that we should prove interesting theorems. And so we try to build a system that is to prove interesting theorems in the predicate calculus. That seems to be ill-defined, but it does also seem to me that I could use an adaptive technique to gradually do better at proving interesting theorems.

A. But that is precisely the difference: The system is the theorem prover, constrained by the rules of the predicate calculus. The system is not the formal structure of the predicate calculus itself. It is the theorem prover that will change state over time. The only source of ill-definedness in your example is then how to measure "interestingness". In other words, an ill-defined criterion on a well-defined system. Also our standards of "interestingness" may change (in ill-defined ways) and thus we have layers of ill- and well-definedness.

B. Returning to your first remarks--why do you want that dichotomy between "us" and "them"? The task of the economic decision-maker is not to control a physical system that knows nothing, but to control a system that is in turn modelling the controllers! And, of course, the economists are part of the economy that is to be controlled. The world, which is the ill-defined system that we are trying to control, is a world full of people. And they are each trying to control certain aspects of a world full of people. And the problem is: How do we coordinate all these different control systems to achieve some sort of overall goal when each control system is changing the environment of other control systems which in turn adapt and evolve. There is no more important problem than that.

A. I don't think it's the same kind of problem at all. Who should define the "overall goal"? Some great social engineer in the sky?

B. No, all of us. The control is distributed. There are indeed systems whose components have different goals, but which can agree on a higher level goal for <u>some</u> kind of overall satisfaction. Perhaps a good marriage is like that. Or a long-lasting peace treaty among nations. So we are all embedded in our own control systems.

A. I may be <u>physically</u> part of the system to be controlled, yet there is <u>logically</u> a separation. Consider my control of my exploratory limb movements. I might analyze this in terms of brain-as-controller sending a command to my limb-as-control-system to move; sensors of the limb then send in signals, and my brain is in a state now to integrate those signals into some building representation of the world. Alternatively, I might just say: I give an order to myself to know more about something--I am controlling myself. The "us" and "them" is a logical separation, but need not necessarily be a physical separation.

B. I would have thought it was still a dangerous separation to make. The word "control" has a certain connotation of brute force about it, which is put in better perspective if we speak of "self-control" or "control of interactions." If you're walking across a street, and there is a fast car racing towards you, one control strategy is to take a machine-gun and try to shoot out the tires so that the car will roll over and avoid you! An alternative strategy is to get out of the way. This suggests the usefulness of studying relationships of reciprocal interaction, like the river bank and the river, rather than of control.

Can we not develop another image of adaptive control which replaces a "brute force" image? Imagine that you have a person sitting down with a very elaborate computer system, of whose limits and total properties he is not sure, and that man and machine are actually trying to work together on a problem. Now, there the idea of making an "us" and "them" distinction looks totally wrong, yet the situation is certainly ill-defined in the sense that the human is not sure what he and the computer are going to do. It seems useful to speak of the cooperative control of a collective of systems, rather than control of one system by another.

A. I gladly renounce any necessity of "brute force" in my analysis. But I would still like to distinguish "us" and "them" in analyzing adaptive interactive control by ill-defined systems. Often, as scientists, we are not trying to control something ourselves but rather are looking at how a collection of systems interacts, and trying to see to what extent they are achieving some global good, as it were. I agree with you that we should not limit analysis to the situation where the "us" is forcing "them" to do something. I accept your idea that we can speak of controlling our interactions--self-control, as well as control of others.

Let me again stress that I think the notion of "ill-defined" only makes sense with respect to some overall task: whether we are stressing an interactive situation, or the task of the scientist as outside observer, or if we do have a task (e.g. spaceship navigation) in which it is appropriate to speak of a dominant controller. I am saying that the systems involved are ill-defined if, in fact, we can't describe them well enough to allow us to meet our control criteria.

B. Surely you're on the wrong track. If a system is deterministic, it is well defined--no matter what the task at hand.

A. I think that the distinction between deterministic and stochastic is totally irrelevant to the question of whether or not a system is ill-defined. Statisticians often feel they have solved a problem when they can use the central limit theorem, yet the central limit theorem is not very interesting for deterministic systems!

B. I can't agree with you. It seems to me that if a system has noise, then you have to say it's ill-defined. Remember, too, that we use the word "system" ambiguously. There is the system as some relevant chunk of the "reality" out there, that we might hope to have some control over; and there is the simplified model or description that will enter into our calculations. If you don't know how the "chunk of reality" works, it may be a deterministic system but you don't know how it works. Your system-model is an ill-defined system until you discover how the "reality" works, and then your system stops being ill-defined.

A. Well, I concede your point to the extent that all we can ever talk about are aspects or models of reality. But I still think determinism is not the essential point here. Consider the coin toss as the classic example. If one is playing a game which involves tossing coins from time to time, one uses a compact description of the coin as a system which on each trial has a .5 probability of coming up heads, or tails. With this system description, one can often develop an excellent strategy for playing the game. Again, treating a well-shuffled deck as non-deterministically described, one can work out winning strategies at, for example, blackjack (twenty-one). One can, I think, say that the system of the dealer and the deck at blackjack is a well-defined system, and one can come up with excellent strategies for winning.

B. Perhaps the mistake is to try to divide the world into two classes. Perhaps the term "ill-defined" system is fine in a conference name, but not as a subject in itself.

A. Groucho Marx said there are two kinds of people--those who divide the world into two kinds of people and those who don't.

B. What I'm getting at harks back to an ancient problem in physics. Suppose I ask the question: "Do you believe that the laws of physics are causal and deterministic?" Nagel answers that: "All the laws of physics are deterministic and causal with respect to the variables they nominate as state variables." But what are those variables? They are the variables that are causal in the equation. So, if you pick your context correctly, physics is entirely causal, including quantum mechanics. It's just that the variables there are probability amplitudes.

A. Yes. In fact, every stochastic system is causal when we describe its state as a probability vector. But, of course, such probabilities can only be measured over an ensemble (unless we can assume an ergodic hypothesis to exploit trajectory averages). Non-determinism is relative to what we can estimate of the state for a given system for a given time. I assume causality <u>whenever</u> I speak of a system, whether deterministic or stochastic--but I say it's well defined when the description is adequate to the task at hand.

B. Ah, I see a different point to be made here. The trouble is that when <u>you</u> specify a task, it is in the language of human intentionality. There is a belief that arises out of too much contact with computers that the world runs by information. But the universe does not run by information; it runs by dynamics that are constrained. And so systems have their lawfulness, and you can't impose human intentionality on them, regardless. That's the trouble here. You <u>can</u> do it inside computers, which are about the worst artefacts to use as images of the real world that I can think of, because switching networks have totally arbitrary dynamics as symbol manipulators. But they are a very poor sample of the real world. And that's the difficulty.

What I'm suggesting is that the problem we're having with economic systems and ill-defined systems is that they are more limiting than we think. It is difficult to control these ill-defined systems because we can't understand the dynamics and the constraints on the dynamics.

A. What you're saying is you can't control it if you don't understand it. Does understanding help?

B. It seems to me that understanding helps because you can play with the boundary conditions on the dynamics. That's the point at which they're sensitive. The paradox here is that the richness of the systems comes precisely out of constraints--a system gets more interesting as it is more constrained. Only up to a point, of course!

However, I don't think "understanding" is equivalent to successful control--which gets me back to my disagreement with your definition of a system as ill-defined if you couldn't control it. In fact, I would argue that a system can be well-defined, but it may well be that you have a model of immense complexity. I have to concede you a point here--I'm prepared now to accept that a model can include stochastic things in it, without being any less well-defined. But if the algorithms required to compute optimal control are too slow relative to the dynamics of the system, then the fact that it's well-defined will not make it possible to control it.

A. Well, we seem to agree, then, that the issues here are complexity issues rather than deterministic vs. stochastic issues. If one has a very large system, one may define it very well by ten aggregated variables and yet know full well that, in fact, there will be random fluctuations about the mean of each of those variables that one couldn't have without disaggregating to millions of variables. But if, in fact, one can show that the variation is small enough that knowing that window of values for the state variables gives one enough information to implement a good control strategy, then the system is sufficiently well-defined for the purpose in hand.

Though we seem to agree on this, I shouldn't let you get away with your slur on computer science. Nonlinear mechanics is fine to describe units in a complex system, but surely the overall structure of such systems--and, even more surely, the appropriate control strategies--must be expressed in some sort of programming-like or procedure-like language. It will probably need concurrency, and knowledge structures of the kind AI is developing, to do it right. It is one thing to "tune" the parameters in a dynamic system to improve performance marginally; it is quite another thing to combine existing structures into wholly new combinations to improve performance by great leaps.

Control theory studies feedback, state estimation, and identification algorithms which can "tune" the parameters of a system description when the system is of a given class of the kind you are talking about. But what if no setting of parameters will do the job? To turn on a slide-projector and to drink from a glass of water requires two entirely different schemas. More generally, interacting with a complex environment will require me to mobilize my knowledge (schemas) about a whole host of objects. To plan our behavior, we need to recognize the world in terms of things with which we can interact, and those things have to be recognized in relationship. Each of those things has to be recognized as a system with certain general properties; and our recognition of it has then to set certain parameters which will allow us to interact with it in the appropriate way. With such an "assemblage," we're leaving the domain of continuum mathematics. Of course, certain aspects of my interaction will depend on mechanics--the way I hold the glass will use things to do with center of gravity and mass. We have the notion of very complex interweaving of computational structures to get access to appropriate control systems to master the complexities of the world around us.

B. Well, be that as it may, I do want to remind you of the point that was illustrated in the discussion of economics: the notion of the other systems themselves having models of you, so that what you do is not, necessarily, simply changing the state of those systems as in physical systems--the state _per se_ as it were (though these are dangerous distinctions!)--but the state of their models. Here the notion of cooperative _control_ becomes very interesting. Oh, I see your point about going beyond physical systems--I suppose you would insist on calling it cooperative _computation_. Borrowing the notion of a _contract_ from law, one interesting idea is that if two interacting systems within this overall control problem will accept mutual constraints, then one may be able to implement a control strategy. Whereas, if those constraints aren't accepted, one may be unable to implement a control strategy simply because one system is going to do what they can to circumvent the other. This whole situation of interaction of actors modelling each other is going to be central to our study of adaptive control, even if we stop using this word "ill-defined." One application on the horizon is air traffic control. The idea is that tactical air traffic control can be distributed into the cockpits, with the ground control stations simply setting the flight patterns.

A. Yes, and it does seem to me that the "contract" is a computer, rather than a control, notion. Anyway, one more remark about the complexity issue. It is extremely hard to conduct useful experiments on large-scale systems, particularly when they involve multitudes of people. How do you even get reasonable data to define those systems or to see whether you are, in fact, accomplishing the objectives of your control algorithms and so on? That's a very, very serious problem that's not faced very often, but is one that will continue to plague those people who are interested in large "somewhat ill-defined" systems.

Another vast point that adds to the complexity of everything is that the evaluation criteria we use often change just as we begin to understand them well enough to use them. Just as I thought I understood that the only important goals were chocolate and sodas, I became a teenager. Similarly, large complex information systems are nearly always used for many purposes not dreamt of by their designers.

B. But surely that is just plain bad design. A good top-down design ought to be flexible enough to allow other usages.

I admit that history is against me. Perhaps that is your point about chocolate: evolution and growth do occur not only in organisms and systems, but also in the tasks assigned to them, and in the purposes and criteria they hold.

A. In that sense, then, if a system is defined with respect to our evolving or changing criterion or set of criteria, it will inevitably become ill-defined it if continues to be useful.

[This "dialogue" sets the stage for what follows. Readers are invited to add their own views to the above--settling some questions, raising others--as they read the chapters that follow. The Editors.]

HUMANS AND THEIR RELATION TO ILL-DEFINED SYSTEMS

Neville Moray

Department of Industrial Engineering
University of Toronto
Toronto, Ontario, Canada

ABSTRACT

Using a taxonomy of systems proposed by Ashby this paper examines the various species of ill-defined systems which may occur in man-machine interaction. A distinction is made between objectively ill-defined systems and effectively ill-defined systems. Various properties of humans as information processors are examined, and it is claimed that humans turn almost all systems, whether initially well-defined or objectively ill-defined into effectively ill-defined systems. It is suggested that while conscious decision making is particularly ill-suited to controlling ill-defined systems, sheer practice frequently enables humans to control them through skills which are not understood even by their owner.

INTRODUCTION

My definition of an ill-defined system is one whose state-transition matrix cannot be known, either because some states are inaccessible, or because some of the transition probabilities are inaccessible, or because the matrix is not time-invariant. Any such system is objectively ill-defined (OID). Although many industrial processes and man-machine systems in general can be to a good approximation regarded as well-defined (WDS), at least as far as the "machine" is concerned, it will be a major contention of this paper that the properties of the human operator are such as to render the man-machine combination of an effectively ill-defined system (EIDS). What is surprising is that nonetheless humans often control EIDS rather effectively.

A NATURAL HISTORY OF ILL-DEFINED SYSTEMS

What features of a system call for it to be classified as an ill-defined system? Following a classification proposed by Ashby (1956) we can divide all systems into three broad phyla, some of whose species are well-, others ill-defined.

State-Determined Systems (SDS)

State-determined systems are not generally thought of as being in any sense ill-defined. They are normally both completely and exactly predictable. All the state variables are known, and the transition matrix is completely deterministic. If the present state vector is known, then all the future of the system is determined and precisely predictable.

Not merely is it difficult to see how such a system could be an OID, it is even difficult to see at first glance what sense can be given to the concept of "control" when applied to an SDS, since all its future is completely determined. However, we may divide the state vector into two sets of elements, one of which is the "machine" and the other the "controller." Let there be two possible trajectories through state space which the system can take, which do not overlap. If the "control" variables are in one state, trajectory A is followed; if they are in the other state, trajectory B is followed. Although the outcome is completely deterministic, the "control" subset of variables has a right to be called a controller. (The question of why the controller takes one state rather than the other is a separate question, and may result in the entire system being non-deterministic. But within the system as defined the concept of control is applicable.) If then a human is controlling an SDS, then providing that the actions of the human are completely determined, the entire man-machine system is still an SDS. In theory at least an operator following exact rules of action laid down by the manufacturer would be such a controller, and a well-defined man-machine SDS should be possible.

Leaving aside for the moment the question of limitations on the determinacy of human actions, is there any sense in which an SDS could ever be ill-defined? One at least suggests itself. In all realizable systems communication takes a finite time. If a system were sufficiently large that the time taken for information to reach the controller (or, more precisely, the observer of the system) were to be long compared with the time scale of state transitions, then it would never be possible to establish the value of the system state vector, and the system would be inherently an OIDS. The controller could never choose the appropriate control state without uncertainty, even though the entire system is an SDS.