ON SYSTEM EFFECTIVENESS

SANJOY K. MITTER[1]

Department of Electrical Engineering and Computer Science and The Laboratory for Information and Decision Systems, Massachussets Institute of Technology, Cambridge, MA 02139 USA

Introduction

To discuss the issue of why certain systems are "effective" and others are not, even though the ambient conditions of two systems may bear close similarities, it is necessary to define the concepts of a system, what it means to be effective and what one means by similarity. None of these concepts are easy to define because they are essentially multidimensional. System effectiveness is intimately tied to the issue of structure, the problem of measurement and the question of resources and values on the basis of which the system is evaluated for effectiveness. I would like to discuss these issues in a somewhat broad context where the systems can include both technological as well as social and economic systems. The central point of this paper will be to argue that, in order for systems to be effective, they have to be coherent in a somewhat technical sense (Rockland, Gaveau & Mitter n.d.). Coherence, however, is also a multidimensional concept and can manifest itself in diverse ways. Nevertheless, it will be argued that systems which lack coherence in certain essential dimensions are not effective. The word *coherence* is being used here in the sense of Whitehead (1978) and is a concept which is broader than logical consistency. It requires viewing the system as a "whole" which always has an environment and a value system (internal). Besides, the system residing in its environment is capable of observation by a multitude of external observers, each observer possessing its own value system.

Thus a system may well look incoherent (ineffective) when viewed externally by an observer, but may well be coherent (internally) when viewed by the system functioning in its environment.

To discuss some of the issues that are being raised, let us look at the Internet as a system. The Internet is a heterogeneous network which serves as a medium of communication for data, images and voice. It has a layered structure, the technological system residing in the lower layers, the applications layer residing in the higher layers, and these layers are separated by the use of two protocols, the transmission control protocol and the internet protocol. The layering structure is more complex, but that is not important for us at the moment. The Internet is a system, it has an environment, namely, the users with their behavioral patterns and possibly other networks such as a wireless network. The Internet is highly effective

[1] This research has been supported by the NSF-KDI Grant ECS-9873451.

in the sense that messages get transmitted correctly to the correct destination most of the time with reasonable delays. It is, however, not efficient. It is, for example, not optimized for minimizing delay. It is a best-effort system which drops packets if a transmission path is not available. It does not guarantee quality of service. We might say that it is coherent. It appears to be "optimized" for robustness against uncertainty, the uncertainty being the uncertain behavior of the users, growth of traffic and failures in the network. The Internet is an evolutionary system and hence, if it is to continue to be effective, it must be adaptive in a changing environment. It also has to be adaptive to changing requirements and values.[2] To give an example, if quality of service becomes an important issue and if pricing mechanisms have to be introduced to ensure quality of service, then it is no longer clear that the current architecture of the Internet with IP and TCP providing the interface between the physical and applications layers is an effective architecture for the service Internet is required to provide. The current architecture is also probably not adequate to provide network security. A system that was coherent loses its coherence in the presence of changing needs and a changing environment. One of the main themes in my talk is to emphasize the dynamical view of systems. The viewpoint of general equilibrium theory in economics with its notion of ideal conditions for operating an economy seems to me to be quite inappropriate—an equilibrium point just does not exist in many situations.

In the remainder of this paper I want to highlight certain conceptual issues arising from recent work in systems, control and information theory which might have a bearing on the understanding of social, economical and technological systems with significant human interaction. The issues I want to discuss are: (a) the dynamical view of systems; (b) the idea of information, its collection and availability in the right place at the right time and the need for setting up communication channels; (c) feedback; (d) the problem of measurement; and, finally, (e) the notion of coherence as an integrative organizational concept.

The idea of a dynamical system

Whenever we talk about a system, we have in mind a quadruple
$S = (S, E, O_I, O_E)$, where S stands for a dynamical system, E its environment,
O_I the observer who resides in the system itself and observes its own interaction with its environment, and O_E, an external observer.

In fact, what we have is a network of such triples interconnected in complicated ways. There are two value systems, an internal value system for the pair (S, O_I) and an external value system attached to the external observer. These two value systems may or may not be consistent or even comparable. We might have a fairly

[2] I use the word *values* in a broad sense which is inclusive of terminology such as performance measures, utility functions in the context of technological systems as well as values in a more philosophical sense.

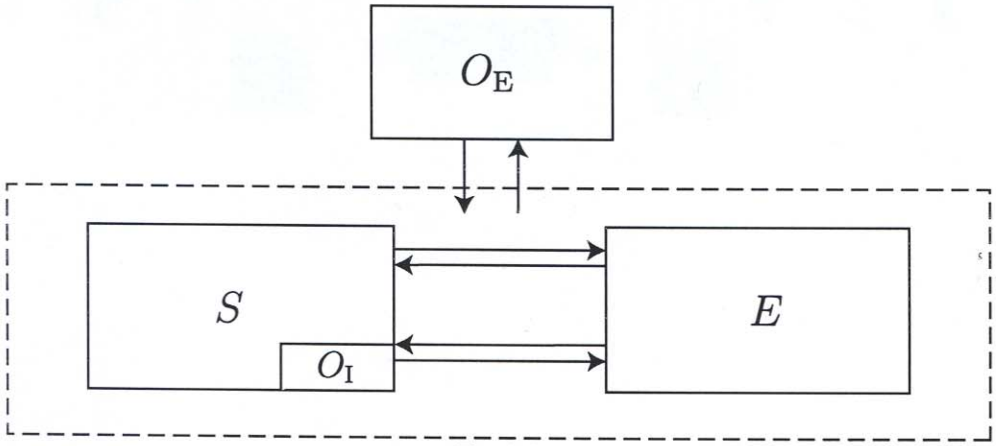

Fig. 2-1:

complete description of the pair (S, O_I) and a fairly incomplete description of the environment. The description (mathematical) is a conceptual scheme to model the system whereby those parts that can be adequately modeled constitute (S, O_I) and those parts that cannot be adequately modeled constitute the environment. Thus, there is a certain arbitrariness to this decomposition.

A system is a pair (U, β) with U a set called the *universum*, its elements are called *outcomes* and ß a subset of U called the *behavior*. The behavior ß is often described by a set of implicit equations (Polderman & Willems 1998).

In many situations the universum consists of manifest and latent variables and the behavior will be a subset of the product set of manifest and latent variables. An example of a latent variable would be price in an economic system. A dynamical system S = (T, W, β) is a triple where T is a time set, W space of signals, and β a subset of W^T. The behavior is thus a space of trajectories. I want to make a point which is important here. Note that a system is not given by a relation between inputs and outputs which imposes a cause-effect structure. Given a behavior, there could be many choices of input-output pairs leading to different representations of the systems with a given behavior. When assessing system effectiveness, these input-output pairs have to be carefully chosen and then experiments have to be designed to measure input-output behavior. It is thus, for example, not at all clear what the definition of productivity should be in a modern economy and, of course, measures to increase productivity and judge its effectiveness depend very much on the choice of input-output pairs being measured. The situation is even more complicated since these choices of input-output pairs have to be made for a system which is under feedback control, a question we shall come back to later. This issue of causality does not seem to have received the attention it deserves in economics and the social sciences. Even in the natural sciences it is a non trivial issue (Weyl 1989).

Communication, control, interconnection and feedback

If a system (S_1, O_1) does not have a desirable behavior, then we might want to change its behavior by interconnecting it to another system (S_2, O_2) called the controller so that the intersection of the two behaviors β_1 and β_2 of systems S_1 and S_2 respectively is a desirable behavior β_d.

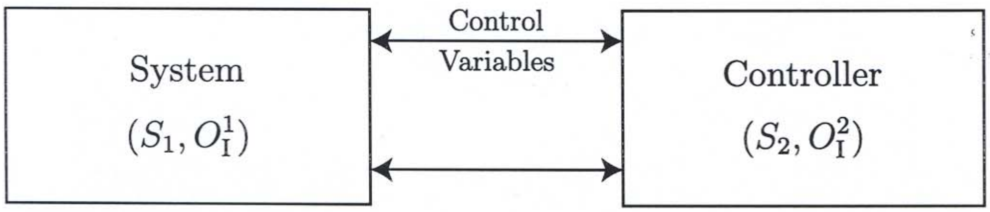

Fig. 2-2

When a cause-effect structure is imposed, that is, inputs and outputs have been chosen, then we often obtain a feedback controller which is diagramaticaly shown below:

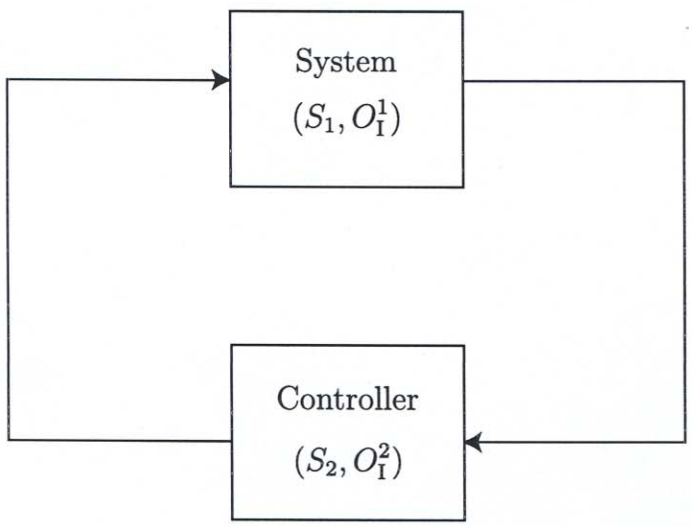

Fig. 2-3

The idea of feedback is a general concept which is quite abstract. For a system to be effective, feedback mechanisms often have to be built in, since the system always operates in an environment which is uncertain and feedback is a general mechanism for reducing uncertainty. The reduction of uncertainty is a key issue in feedback control. In feedback control one has to decide what needs to be measured and what signals have to be "fed" back from the controller to the system being controlled. Thus feedback control is intimately tied to the question of measurement and implementation of control actions.

In a distributed environment such as that of an organization which consists of a network of system modules as described earlier, feedback will be affected not on physical signals but on information extracted from physical signals. Thus physical signals have to be coded into information and since this information will be distributed throughout the organization, communication channels will have to be constructed so that this information is available at the right locations at the right time. I believe that the absence of such communication channels is one of the primary reasons many organizational structures are ineffective. Without the existence of such communication channels, effective feedback contol cannot take place. Indeed, feedback itself may be viewed as a communication channel, and this communication channel has to be created and implemented. The role of coding and decoding in these situations is a difficult subject and will have to be discussed elsewhere. However, the availability of coding and decoding possibilities helps to define and perhaps fundamentally change the architecture and organization of a system.[3]

I should add a note here. When I talk about communication channels, I mean both communication in a technological sense and human communication, both individual and collective. It would be interesting to examine the ideas of Habermas on the theory of Communicative Action in light of the sciences of systems, communications and control. It is unclear how dialogue and communicative action in the sense of Habermas can be organized and implemented (McCarthy 1978).

Values: internal and external[4]

I have so far given you a sense of how a systems theorist thinks about a system, how it is situated in an environment and how it is observed by the system itself and also by an external observer. I have emphasized the need for creating communication channels so that information is available at the right place at the right time and the importance of dialogue and communication within a network of systems. All systems will operate in an uncertain environment and, if a system is to be effective, it must function robustly and this necessitates the need for feedback control. This feedback control may change the signals entering the system but also may change the structure and organization of a system. In order to be adaptive, the system must have a certain degree of autonomy and this must be built in or the system must be capable of attaining it. There might be lessons in this for the global economy and the way it functions. Local autonomy through political and human action is necessary for the global economy to be robust in the presence of uncertainties.

[3] For an important discussion of some of these ideas in the context of economics, see Arrow (1974).

[4] This section is taken almost entirely from *Autonomy and Adaptiveness: a perspective on integrative neural architecture*, CICS Report P-431, MIT, September 1994.

Suppose we have made provisions for the various elements, like communication channels, coding, decoding and also structured it for autonomy, adaptiveness via feedback and possibly other forms of control. How do we actually design these systems? How do we evaluate these systems if they are autonomous (natural and social systems)?

Goals, values and optimality

(a) Goals

An autonomous system has its own goals. By this I mean that its goals are not fixed by a designer, but are emergent properties of the system in its interaction with its changing environment.[5] By "own goals," I also mean that the complex forms of behavior are not *a priori* pre-specified and "pre-wired" but that they are emergent. In particular, there is not a fixed repertoire of behaviors, organized into a pre-set priority hierarchy. The term "goal" as used here is very broad. For example, it need not be tied to any "awareness" "or consciousness" or "volition" on the part of the system. Nor need it convey any intentionality insofar as this is directed towards any specific object or aim or end external to the system. The minimal goal (or proto-goal) is implicit and system-referred:[6] to maintain itself as a system. This requires maintaining a state of relative adaptation with the environment and, in particular, maintaining metabolism within safe bounds. Outward-directed goals emerge as conditional concomitants of the system's modes of fulfilling (and of giving explicit form to) its implicit proto-goal, and these modes are themselves conditioned upon the systems' internal organization and "resources," and on the state of the environment.

"Emergent" does not mean emergent out of nothing: there must be a kernel or ground starting from which the system in interaction with its environment is able to "bootstrap" itself. This kernel includes a physically embedded structural/functional substrate which supports mechanisms to induce and sustain the fulfillment of the proto-goal. We shall, somewhat metaphorically, refer to this structural/functional substrate as the "basic circuitry."

[5] This characteristic of "autonomy" suggests that my use of the term is different from that of designers who seek to build "autonomous robots." That is, is it possible to arrange the design so that the system is autonomous in our sense, yet at the same time meets specific goals of the designer (e.g., specific tasks to be performed, or ends to be achieved in the environment). Perhaps this is to some extent realizable in the form of additional proto-goals built into the "circuitry" and "value systems." However, it is unlikely that a specific task, such as driving a nail into a wall at a specified location on the wall can be "emergent." More likely what is desired is a "quasi-autonomous" system with the following characteristics: (i) it can act autonomously (i.e. "take care of itself" and "not cause damage") when not required to perform specific tasks, (ii) when required, it accepts and is able to interpret and execute (in a context-dependent fashion) instructions (expressed in a high-level language) specifying the task to be performed but not the details of execution.

[6] Making a rather strained analogy with a formal system setting, the idea of system referral vs. external-referral, is akin to a recursively-defined condition on a function or process.

A couple of points of clarification: (i) the "basic circuitry" (certainly the particular form/function complex comprising it" is not "canonical," and may vary from system to system. Moreover, the basic circuitry in any individual system need not comprise the entire "circuitry" of the system, and must be integrated with the rest of the system organization. In particular, the "basic circuitry" should not be confused with the "initial circuitry" of the system, that latter being the initial state (at "birth" or time of "construction") of the overall "circuitry" of the system; (ii) the circuitry and value system are not static. Over time (i.e. the "somatic time" of the system), both the "circuitry" and the value system, in interaction, will undergo modification, and may or may not reach a quasi-steady state, or pass through a succession of such states. In particular, the value system, which initially consists only of the primary value system, will expand to incorporate "higher-order" values, as well as "symbolic" values, which serve to enhance the adaptiveness of the system, involve processes of learning and memory. To an extent, the substrates supporting these processes may themselves be emergent, and undergo modification over somatic time. Referring to the "time scale" terminology discussed below, I shall (essentially as a tautology) assume that the "initial circuitry" and "initial value system" of the given system are the result of evolution (resp. the designer) working over evolutionary time (resp. design time). The other constituents of the value system will be emergent. Here, and elsewhere, when I speak of "emergence," I shall mean emergent in somatic vs. design time, unless I explicitly state otherwise.

(b) Time axes

I wish to allow significant structural changes during the entire lifetime of the system. Thus, I shall not be concerned to distinguish developmental time, and shall generally subsume it under somatic time. Moreover, in the context of artificial systems, I shall susbtitute "design time" for "evolutionary time." This is the time during which the designer may generate and test a (parametrized) family or "population" of designs. I picture this design time not so much as a distinct time *scale*, but as a distinct time axis, interleaved with the somatic time axis, During design time, any testing, modification, or selection is one by the designer on a population of systems or of system designs. During somatic time, the designer does not intervene, and there is only the individual system and not a population. The only "exploration" is that done by the system, and the domain explored is the environment, via direct interaction, rather than a "design space". Such interaction may result in (beneficial or harmful) modification of the system, including modification of the "circuitry" and of the "value system." Such modification may, in particular, take the form of "learning." If the system has the capacity for "planning," this provides in somatic time, a kind of "virtual" population, as a sort of partial "surrogate" for the actual population of design time.

(c) Basic circuitry

I imagine the "basic circuitry" would, at a minimum, contain mechanisms which manifest sensitivity to alterations in the internal state of the system, as well as mechanisms capable of restoring the internal state to a "baseline" or "dynamic eqilibrium" level. This may require more than a simple feedback loop, e.g., it may require significant internal structural or functional reorganization of the system achieved *via* internal effectors acting "mechanically" in conjunction with cascades of "chemical" processes. Moreover, the "mechanical" processe may themselves need the support of the "chemical" processes to fuel their activity. In addition, the activity of one set of effectors or processes may give rise to the need for additional mechanical and/or chemical activity, for example, to eliminate waste products. Part of the "job" of the effectors may also be to alter the situation of the system in the external environment, in particular, *via* locomotion. The system will also require "interfaces" with the external environment in order to take in resources and to eliminate waste products. I do not expect the interface to be a purely passive conduit. Hence, there will be the need for effectors to physically control the activity of the interface, as well as mechanisms sensitive to alterations in the state of the interface. The state of the interface should undergo some (localized or diffuse) alterations induced by physical stimuli impinging on it from the external environment, so that the system has the possibility of taking account of "resources" or "risks" in the environment. (I expect the basic circuitry to act predominantly in "automatic" fashion, i.e., independently of direct control intervention by other parts of the circuitry. Byt this I mean that the remainder of the system does not directly monitor, and hence does not directly regulate the state of the basic circuitry; rather, it influences the behavior of the basci circuitry only by influencing the state of the external environment of the basic circuitry.)

I have spoken of "sensitivity" rather than of "sensors" so as not to appear to suggest some *a priori* modes of categorization by the system of the sensory "inputs." We expect this categorization to be an emergent process. Similarly, the sensitivity may be of a diffuse vs. sharply localized or sharply specialized character. The system need not be sensitive to precisely where a signal has impinged, nor need the system cleanly segregate signals of different physico-chemical type. Nor need the signals result in elaborate internal representations of the environment; for example, photosensitivity need not imply a sense of vision. The sensory and effector apparatus may become refined and elaborated over time (evolutionary/developmental/ somatic).[7] This refinement will typically go along with an elaboration and reorganization of the overall structural/functional architecture of the system. The sensory apparatus need not function in a purely "passive" manner. Rather, the system may actively "explore" for "relevant signals in the external or internal environment.

It is not conceivable that, from the standpoint of a strictly formal/logical vs. embedded system, a complete list of the logical "components" required for (one of the variants of) a "basic circuitry" may be very brief (as in the case, say, with von Neumann's "self-reproducing automata"). However, when one considers an actual

[7] I do not rule out malfunctionality. Thus, the same physics component or activity pattern may serve both sensory and effector functions.

embedded circuitry, the situation may be quite different. As hinted at in the above partial enumeration, the inclusion of any one component gives rise to requirements for additional supporting substrates or processes, which in turn generate support requirements of their own, etc. By the time one attains "closure" or "self-containment" (in dynamic equilibrium with environmental resources), one may have a very complex system. For example, a cell, even if we disregard apparatus specificaly associated with replication (the cell cycle), is tremendously complex. Nor would we say that a virus (which draws on a host cell's reproductive apparatus) is "simple".

(d) Value system

An adaptive system, embedded in a complex environment, cannot try out actions completely "at random" all the time. (Here "at random" may, but need not, connote "stochasticity". The emphasis, rather, is on "non-directedness".) Even if certain actions are done at random, as may particularly be expected early on in the "lifetime" of the system, there should be a system of values or rewards which has the effect of retaining "valuable actions" for the system. I envision this value system as an outcome-based biasing system, rather than as an encoding or representation (explicit or implicit) of pre-set behaviors, indeed, in general, the value system will have no way of "knowing" (nor need there in principle be a way of knowing) what the (*a posteriori*) beneficial behaviors or goals actually are. In particular, the value system should be consonant with situatedness, or context-dependence, of action, and with behavioral choice. As discussed above, initially the value system consists exclusively of the primary value system, which is associated with the minimal proto-goal, and is realized via the basic circuitry. Over the course of (somatic) time, the value system can expand, with certain new values appearing, and others disappearing or being modified, depending also on changes in the environment. As emphasized above, these values are emergent, and not imposed by an external agent.

In speaking here of a "biasing system," I shall not seek to distinguish between the system for generating the "biases" themselves which supports the "biases" thus generated, or the "biases" themselves. (For that matter, these distinctions, certainly that between a bias and its substrate, are quite murky.) In particular, the primary value system involves all of the above, including certain "pre-set" biases determined in evolutionary time (or in "design" time). Nor do I mean to suggest a uniformity or homogeneity of "mechanisms". As one example, even in the case of the primary value system, certain of the primary values, such as those directly linked to homeostatic constraints, may involve error-correcting mechanisms of traditional "feedback" variety. Others, probably constituting the general case, will have more of the "unspecific" or "diffuse" character emphasized in the preceding paragraph. I should also note that when we speak of individual "biases" or "values" as entities in their own right, we do not intend this in any literal sense, first, because a particular "value" makes sense only as part of an overall system, second, because, as emphasized above, values in the sense we have in mind, tend to manifest a kind of generalized or diffuse vs hard-edged quality.

(e) Optimality

I believe that the "optimality metaphor" does not apply to the workings of an adaptive system. In general ad adaptive system does not try, either through its own "intent" or through the nature of its underlying "design", to optimize anything in particular. There may be conflicting, incommensurable criteria.[8] When the system, or some subsystem, does optimize a particula "utility function," this utility criterion may readily vary from one context to another. Moreover, even when a behavior does in fact "optimize" some criterion, we expect this criterion itself to be emergent from the interaction of the system with its environment and not, in general, prefigured in the "design" of the system.

The need for control: achieving and maintaining coherence

All adaptive systems (whether natural or artificial) will tend to be rather complex. Even to maintain themselves as systems, let alone to accomplish a specific set of outward-directed goals or tasks, they will need to integrate the activity of multiple interactive subsystems. Moreover, this internal organization of structural resources into functional subsystems may be dynamic rather than static. In addition, distinct "levels" of subsystem "decomposition" may be superimposed. (This is illustrated most readily in the case of societal organization: among the superimposed levels are the political, legal, economic, familial, etc.)

However, a collection of subsystems put together willy-nilly cannot be expected to yield "coherent" behavior, let alone behavior continually maintaining a state of adaptation with the environment. Even if the subsystems are themselves individually "coherent", they may very well tend to compete in an anarchical way for access to various resources. (This need not be limited to a societal organization: to take a very simple biological example, antagonistic muscles could attempt to work simultaneously.)

One is thus faced with the question of "control": namely, what are the means by which the system achieves and maintains coherence in a context of emergent functionality? I feel that this conception of the problem of control will require that one goes well beyond the traditional perspectives of control engineering. Even in the setting of artificial systems, the various subsystems may each have *several* naturally associated "utility functions." Correspondingly, the design of the control

[8] I think that the "optimality metaphor" is fundamentally out of place in this setting. In particular we don't think that any aid or comfort is to be drawn from the techniques of so-called "multi-criteria optimization," e.g., forming new utility functions from weighted averages or component utility functions, or drawing on Pareto optimization ideas from microeconomic theory. We note in passing that in Pareto optimization theory, one does not have a canonical way to choose from among the many Pareto optima.

system may involve complex trade-offs.[9] In particular, it will be necessary to reexamine the extent to which control can maintain its traditional "supervisory" character. In a setting of emergence and situatedness, a putative supervisor (associated with any aspect of the overall organization) might, in principle, have no way of "knowing" (certainly not to a high degree of "specicifity") what task, state-of-affairs, internal state, etc., should be definitely achieved or definitely avoided by that portion of the system under its influence, let alone what behavior should be "induced" so as to bring this about.

Modes of failure and the problem of measurement

There are a number of related issues which I will not be able to go into in detail in this paper. One is the question of specifications and the "polar" notion of constraints of the system. I want to emphasize that the specifications for adaptive, autonomous systems need to be vague, that is, not rigidly defined. Furthermore, specifications cannot be just declared but they need to be put into effect, perhaps in the form of coherent sub-specifications.

It is important to understand the various modes of failure of adaptive autonomous systems. These failures may be failures to meet specifications, incoherent evaluation criteria, inability to anticipate contingencies, overspecification and dealing with the possibility of system failures.

For adaptive autonomous systems which are partially designed and partially autonomous, mechanisms of measurement have to be built in so that the dynamic behavior of the system can be monitored. Simple regression analysis or more sophisticated statistical testing are not enough. The ideas of dynamical measurement through an appropriate sensory apparatus, the resulting processing of that information with a view to feedback control is an idea which has played an important role in control systems theory. A suitable reinterpretation of these ideas for the structuring and evaluation of socio-economic systems would be a fruitful area of research.

Conclusion

In this paper I have tried to outline how a systems theorist would view the problem of effectiveness in the context of engineering, social and economic systems. I

[9] A separate question which we shall not take up here, and which is perhaps relevant primarily in the design (i.e., artifical system) setting, is that of incrementality of control design. For example, one may not want, when "gluing together" susbsystems, to have to do substantial internal "rewiring" of the subsystems. In that case, one must constrain the control design so as not to permit (or require) it to make use of detailed knowledge of what is happening internally within the subsystems. On the question of multiple utility functions and their coherence see Changeux & Connes (1989).

believe that a dialog between economists, social scientists and systems theorists can bring about a synthesis of views which would benefit economics and social science.

Acknowledgements—The ideas presented here represent joint work with Bernard Gaveau (University of Paris VI) and Charles Rockland (University of Iowa).

References

Arrow, K. J.
1974 *The Limits of Organization*. New York/London: W. W. Norton & Co.
Berthoz, A.
1997 *Le sens du mouvement*. Paris: Éditions Odile Jacob.
Changeux, J.-P., & A. Connes
1989 *Matière à pensée*. Paris: Éditions Odile Jacob.
Damasio, A.
1994 *Descartes Error*. New York: Grosset-Putnam.
Edelman, G.
1987 *Neural Darwinism*. New York: Basic Books.
Gaveau, B., C. Rockland & S. K. Mitter
1994 *Autonomy and Adaptiveness: a perspective on integrative nural architecture*. (CICS Report P-431). M.I.T., Cambridge, MA.
McCarthy, T.
1978 *The Critical Theory of Jürgen Habermas*. Cambridge, MA: M.I.T. Press.
Polderman, J. W., & J. C. Willems
1998 *Introduction fo Mathematical Systems Theory*. Berlin/New York: Springer Verlag.
Rockland, C., B. Gaveau & S. K. Mitter
n.d. *Autonomy and Adaptativeness*. Cambridge, MA: M.I.T. Press. (Forthcoming).
Schönberg, A.
1994 *Coherence, Counterpoint, Instrumentation, Instruction in Form*. Edited with an Introduction by S. Neff and translated by C. H. Cross & S. Neff. Lincoln, NE/London: University of Nebraska Press.
Weyl, H.
1989 *The Open World*. Woodbridge, CT: Ox Bow Press.
Whitehead, A. N.
1978 *Process and Reality*. New York: The Free Press.

Discussion

AUSENDA: Can one define the efficiency pertaining to an Internet system?

MITTER: Well, there can be many definitions of efficiency (as distinct from effectiveness). You can say, well we define Internet to be efficient if, for example, its error rate is below some threshold. Another definition of efficiency might be the fact that it delivers packets of mail on the average with minimum delay. Another possible definition of efficiency might be maximizing throughput through the network. So, I don't think that there is a single definition of efficiency that one can think of, but you might say there are potentially many definitions of efficiency. So, you know, some conplex combination of these various performance measures might lead to a global definition of efficiency.

AUSENDA: This is clear, but does energy utilization have any importance in subsystems?

MITTER: I don't see that energy enters into the picture in the above discussion. There doesn't seem to be a relation. There might well be a limited amount of energy available.

AUSENDA: Energy expenditure in this case is negligible.

MITTER: [Answering a question on the nature of the effectiveness of the Internet and not efficient] All I was saying was that the Internet is effective in the sense that most of the time, it delivers packets or mail or whatever in reasonable time delays. And that for very very large numbers of users its performance is quite satisfactory. So, in that sense, I was using the word effective.

ARONSON: Would you clarify please what you mean by "robustness"?

MITTER: What I mean by "robustness" is, for example, that the Internet seems to work well for a hundred users or one million users, it seems to work well in both cases, and hence it is "robust" against the number of users. You could have other measures of robustness. You could make probabilistic models about the traffic on the network and you could ask if these probability definitions are perturbed somewhat, does the network still accommodate satisfactorily these changing hypotheses? So that is another measure of robustness.

AUSENDA: So, in essence you are saying that there is no possibility to measure effectiveness. It is just a quality which says that a system is more or less effective, without being able to measure it.

MITTER: I am not quite saying that. I think that you could make hypotheses about what constitutes effectiveness, and then you could try to test them out. So, let's make the hypothesis that the Internet is somehow optimized for robustness against uncertainty. Now you can devise experiments in order to try to test that hypothesis. To me, measurement of effectiveness can't be reduced to some utility function, or something like that. It is a complex of things.

MESSERI: Is "coherence" a variable property, or a discrete property? Are certain systems coherent and others not? Is "coherence" a broad general cover all for ordering?

MITTER: Let me give you a kind of a mathematical example. De Finetti was interested in subjective probability. You have "events" or entities, and an entity was assigned a number between 0 and 1. So you had an entity and an assignment of numbers between 0 and 1. And he has a definition of what a coherent assignment of probabilities is—probabilities is not quite the right word. It turns out, that the definition of De Finetti of coherence in this finite situation is completely consistent with the existence of probabilities. You can prove De Finetti's system of entities, let me call them "events," and his assignment of probabilities, let me call them probabilistic assignments are coherent and if and only if the algebra generated by the entities has a probabilistic structure.

ARONSON: Can you explain what "coherence" is?

MITTER: "Coherence" is a very broad concept and it has very many manifestations. For instance, if I listen to a certain kind of classical music, that's coherent to me, so that's one level. At the other level there is logical consistency. You can test if a certain proposition is true or false. That would be the other level. I do not want to give a rigid definition of coherence because I don't think it is the correct thing to do.

THOMPSON: Going to the next question on chapter 4, if you can't give a definition of what "coherence" is, then how can you make a statement saying that the collection of subsystems put together willy-nilly cannot be expected to yield coherence? And you put coherence in parenthesis. But I would say that the Internet for instance, to give an example, does in fact adapt.

MITTER: I gave you a probabilistic assignment of being coherent. I gave you another one, namely of a logical system, and you have propositions that you can test the truth or falseness in that logical system. So, consistency is a kind of coherence, an extreme form of coherence.

DAUN: I read "coherence" as meaning value consensus in the interconnecting structure.

MITTER: Yes, but again it depends on the perspective, internal or external. You may have coherent systems from the point of view of the system itself, but looked at externally they may not look coherent to you. So you have to look out for that. In my view, for complex systems I do not think you want to make a rigid definition. Because, if you do, you are going to get into trouble. But, you can work out the consequences of coherence. You can make a definition and you can work out the consequences of it. It is in that sense that I am using coherence as a conceptual paradigm.

AUSENDA: I have a question on page 3. You define a "dynamical system as being endowed with an internal and external value system." Could you explain what these are, and what the differences between the internal and external one are when applied, for instance, to Internet?

MITTER: I don't think that anybody knows how the Internet actually operates. If you ask me to describe what the laws of operation of the Internet are, I wouldn't be able to do it. But, if you are able to imagine for a moment that you *are* the Internet that, in some sense, you are able to enter the internal circuitry or to monitor the algorithms that are going on within the Internet itself, then one could speculate that the Internet appears, given its resources, given its environment, to operate coherently. But for an external agent, if your criterion for measuring how the Internet should operate is some external criterion, like it must receive packets within a certain time element, then it may well appear to you from the outside that it is incoherent. So, that's what I mean by the internal and external definition of "coheremce."

AUSENDA: As to advantages, is coherence the same as effectiveness?

THOMPSON: If you lived in Yemen and your telephone line was out everything would go on anyhow.

MITTER: I haven't even touched the whole issue of culture in this. Let me say one thing: what we found was that the appropriate language for describing a complex system doesn't seem to exist. Leave alone the conceptual foundations or laws governing a complex system. I am not proposing a kind of universal theory in any sense. I don't think that is even a good idea, but you can still look for certain conceptual themes, broadly defined, and work out its consequences, and coherence is one of them. I am trying to make a distinction between a simulation of a system and understanding its laws. When you simulate something, it is not at all clear that you are discovering the laws of the system or anything of that sort. So, the solar system can be certainly described by Newton's laws, you can certainly model it by Newton's laws, but nobody can say that the solar system is solving a differential equation or anything like that. You have to tell me what you mean by modelling.

MITTER: [On scaling down models of systems that do not have the same macro-level behavior.] There are also reasons to believe that there is a phase transition in systems like the Internet. So, as in statistical mechanical systems, there are things like critical quantities, critical temperatures, and the behavior is of one kind above this critical temperature and below it is completely different.

THOMPSON: I think that we are going to have a problem with the concept of scale. And certainly in my field, in development aid, what happens on a micro-scale is totally different from what happens at the macro-scale when the same program replicates. For instance, in immunization programs, the international programs tend to do completely different things.

MITTER: I raised the issue of time scale, but heterogeneity is another big issue. Page 4.

MITTER: [On input, output and causation and on the fact that when there is feedback the issue of causality is blurred because it could be attributed to both ends of the feedback network.] If you speak about an economic system, it is certainly being controlled. When I say controlled, I mean through government action, money supply, etc. It is not a system that is operating "open loop," but there is a sort of a feedback round the system. So, what you mean by causality in a feedback systen, is a very difficult issue.

AUSENDA: On the issue of "causality [that] does not seem to have received the attention it deserves in economic and social sciences. Even in the natural sciences it is a non trivial issue." Could you give an example of the fact that it has not received the attention it deserves by economists?

MITTER: Well, I am not an economist, but let me talk about models, as when you try to build an economic model based on data. As far as I can see, the framework is more or less an input-output framework. There has been very little work on modelling a system that is under feedback control. So, if you believe that feedback reduces complexity, the action of control actually changes the system and its behavior. So the question is, "what can you identify or what can you get at, from data, which is being collected on a system which is being controlled?" That is a difficult issue. And to the best of my knowledge that sort of issue has not been touched in economics and social sciences.

POZZO: Actually the reliability of the source is an issue in economics. So, in macro-economics they deal with sources and they make beautiful models of them, but you know that you can't control them. For example, the Thai and the Korean government tend to release figures on which one must put a question mark.

MITTER: But in econometric models from data which you estimate per annum, it is not clear that during the control period you can actually estimate all the parameters—the model may be over parametrized. There is a whole set of issues related to this. Whether unemployment is a cause or an effect depends on the theoretical construct you have in mind. So, you can have theories and you can try to test them. But inputs and outputs are not given to you a priori.

THOMPSON: That I accept, but I think that the idea that economies can be controlled is no more than an idea.

MITTER: Economic systems are being controlled, for example by varying the money supply.

AUSENDA: I believe that the economic system controls the money supply; it is not the other way around.

THOMPSON: Take an extreme form of control, embargos as in Iraq, or Rhodesia or in Vietnam. They don't work. They have never been effective. And that is a very extreme form of control.

MITTER: No, I do not say all control actions are going to be effective. What I am saying is that economic systems are being subjected to control actions.

AUSENDA: I think that you are pushing too far your simile with circuitry. I do not think you can have feedback control on such a complicated system as an economic one. I think that all Greenspan does is that when he sees that the money supply is decreasing, then he decreases the money supply.

MITTER: Now, wait a minute, I am not specifying for you what all possible forms of control of an economy might be. There may be inefficient control mechanisms in operation which I am certainly not able to specify.

VOICE: You can view it in different ways and say that the economy controls Greenspan.

MITTER: Sure.

MESSERI: You don't like control in economics.

MITTER: In economics, it may or may not make sense.

THOMPSON: The Japanese could not control their banking system during the last twenty years.

MITTER: In control, the concept of "controllability," namely your ability to transfer the state of the system from one level to another is a well defined concept. Now, a system may or may not be controllable. And this depends on the choice of the control variable, or what you measure and so on. So, the fact that you take control action doesn't mean that it accomplishes exactly what you are trying to do. That's not what I am saying.

VAN WILLIGEN: One aspect that I suppose we could say is cultural is that there are publicly stated actions which may be ineffective for achieving their goal, but then the interpretation of their effects takes place. The fact that control systems exist, they may not be effective or their effectiveness may be limited, but the meaning of their application has a multiplying effect in the system, and so people begin to act in anticipation of their effectiveness.

MITTER: Let me give you a kind of an impossibility theorem. You are trying to control a system. On the system there are all kinds of external disturbances so, for example, the demand may change in some random fashion and you do not know how it changes. Now, what you want to do is you want to regulate a system against the bad effect of these disturbances, and at the same time maintain stability. You do not want the system to be unstable. Now, in very limited mathematical settings, one can prove that you will only be able to do that unless you have a rather reasonably good model of the disturbance itself. Otherwise you will not be able to do it. So that is a kind of an impossibility result, a fundamental limitation kind of theory.

DAUN: If we look at communication, which is something going on in a system, it can be very different: we have anything from cybernetics, to the most simple structure of exchange of feelings, emotions or whatever. In education it is not the whole personality that is involved, but in the living world they are almost always involved in a dialogue all the time. So, the question is, "could communication in a living system be in contradiction with respect to communication in an artificial one?" As to communication in educational systems: it could consist of an exchange of simple messages from children from the age of 5 up to 11 or 12. This has to be a simple communication of some kind, but maybe the system is characterized by all these features in the complexity that Habermas mentioned. But when we go to I.T. (Information Technology) where I am supervising a Ph.D. study via e-mail, the student made a qualitative study based on some interviews, and we were able to exchange messages between Sweden and Zimbabwe to a certain, what I call, "superficial" level. However, when it comes to deeper things, it was impossible to discuss the interpretations of what she had found: I think there are limitations to I.T. When you talk about efficiency in education you have to distinguish between the different types of communication which apply in each case.

Page 8

AUSENDA: By emergent behaviors which are not pre-wired, do you mean that reactive behaviors can be almost random and the environment would in some way privilege the most appropriate?

MITTER: I do not know. They are not random, I think.

AUSENDA: If I understood correctly, middle of the page, "The primary value of a system is that of sustaining the fulfillment of the proto-goal, i.e., to maintain itself as a system"?

MITTER: In the second paragraph I wrote that in this context memory, learning and maybe selection all play a part. What I am trying to get at is that these goals, values, are in general implicit and cannot be explicitly specified. One simply has to deal with that. I do not know how to deal with them in a technical way, but maybe one can discuss them technically in a very limited situation.

AUSENDA: On page 8: how are those higher order symbolic values acquired except by external conditioning as a consequence of the reaction of the environment to the random behavior of the system?

MITTER: Well, they are acquired through the interaction of the system with the environment, I will agree with that.

AUSENDA: But I think that it is the random behavior of the system, not that of the environment. In other words, the environment reacts to the random behavior of the system, and the system acquires that experience.

MITTER: As I see it, the system is in some sense conducting experiments about the environment.

AUSENDA: For instance, when we talk about evolution, is it sometimes defined in a way that the system impinges on the environment and is making experiments with the environment?

MITTER: I think so. So let me give you an example. Nearly fortyfive years ago, my colleague Jerry Lattvin wrote a paper, called "What the frog's eye tells the frog's brain." And the conclusion they came to is that, the nature in which information about the environment is transmitted from the frog's vision system to its brain, is highly abstract, it is not like, "I look at this wall, and I look at every point of light intensity on the wall, say. But they see it in a highly abstract form." For example, it looks at, perhaps, the boundaries of the objects, it makes a representation of that boundary, so that's the sort of symbolic description that seems to be going on. That's one example.

ARONSON: It is no less true of the human eye, the human vision.

MITTER: Yes. The question is how does one make a theory of that? What constitutes recognizing the door as a door? There is a syntactic part to it and also a semantic part—what might a theory of that be? So the concrete work that we are doing is related to that subject. So, when I say symbolic vision, I mean something like what Jerry Lettvin was alluding to about the way the frog's eye communicates with the frog's brain.

ARONSON: By any chance do you deal with animals rather than humans?

MITTER: Not necessarily.

ARONSON: Because it occurs to me from discussions around the day that you are exclusively focussed on human events. Is that true?

MITTER: Yes, that is partly true, I started with a system we are talking about, in a technological context, such as the Internet. On the other hand, there are human interactions going on over the Internet.

ARONSON: OK, for me the Internet qualifies as a human system, to the extent that, when you talk about imput and output or throughput, there is no consideration of what happens to the unit of energy or other product inside the system. Something comes out, it may be the same sort of thing as went in, but may also be entirely different. The evolution of the trajectory inside the system, I think, is of critical importance.

MITTER: Yes.

ARONSON: As a biologist I care, because the systems I am interested in have different habitats and different internal processes as well.

MITTER: Let me clarify. From the system theory point of view, you can look at a system, let me call it "externally." There are two sets of variables, and relations between them, and you can view the system externally. Now, we also need to know what's inside the system. So, in a technological sense, we may want to know how the system is actually implemented. So we would call internal states or latent variables the states of the mechanisms inside the system.

ARONSON: I like "internal states" better.

MITTER: "Internal states" is indeed a better word. I also give a definition of a state. It relates to the internal behavior of the system, excepting in the case of an Internet, I don't know what the internal states are. I mean, it may be a description of the actual algorithms that are taking place inside the system.

ARONSON: My point is very simple. I think that what happens inside the system to the unit in question, is a process rather than a state, and I think that it should be included in our reflections and, possibly, in our definition of "effectiveness."

MITTER:Yes, when I talk about a dynamical system, I just do not mean an external description, I also mean the internal description of the system, if you like. This is described as the evolution of the state of the system, what you call process. Page 9.

AUSENDA: One third of the page down. When you substitute "design time" for "evolutionary time," you define it as the "time during which the designer generates a parametrized family of designs." Does the fact that the designs are not randomly produced make the "design time" shorter than "evolutionary time"?

MITTER: Well, what I am trying to get at is that there are time-scales involved in design.

AUSENDA: I am comparing two designs, on one side I am comparing a "willful design," where things are tested according to certain given projects and goals, and on the other "non-willful designs" such as take place in nature. And so, my question is—maybe the people who work with trees and animals can answer it—does that mean that goal-directed designs achieves a goal faster than evolutionary designs which are random?

MITTER: When I am talking about "design time" here, I am talking about artificial systems, I am talking about the design of an aeroplane, say, and there may be several examples, and anyone could try to understand, to compare those designs.

AUSENDA: What I am trying to find out is whether an artificial system has a shorter design time than a natural system.

ARONSON: If I may make one comment: in evolutionary systems, there is a great deal of randomness.

AUSENDA: That is exactly it.

MITTER: It is not at all clear what you mean by "randomness." When you raised that issue, I mean are you calling it random because you do not know how it functions? What do you mean by random?

ARONSON: In the context of natural selection in evolution, people make the kind of statement that Giorgio [Ausenda] seems to share, set it up in opposition to the creationist argument, or God or some spirit or force as designer of natural systems. That natural selection operates faced with a wide range of—I am trying to avoid using the word "random"—but something like random arising in the emergence of phenotypes and a selection is made.

MITTER: But I do not know. I wouldn't use "emergent" and "randomness" to be the same thing.

ARONSON: Nor would I. I was just playing with the word "emergent" because I really don't understand your use of it. So I wanted to bring it out on the table.

MITTER: My use is in the sense that where it is highly implicitly defined there may be several loops in action which you cannot explicitly specify. But with interaction somehow robust, adaptive systems could emerge. That is the sense in which I use the term.

ARONSON: In biology, there is the term "epigenetic." Is that similar?

MITTER: Yes, it is similar.

Page 10.

MESSERI: My first question is, do you give a distinct definition of values?

MITTER: I am using values in a sense which is quite general.

MESSERI: I think a value is generalized orientation which tends to be a construct for explaining why certain kinds of behaviors are taken over by others within the situational construct. And I also agree with you that a value cannot be bound up in the situation. They have to be added to it or they are situationly bounded.

POZZO: Values are the conditions and the possibility of a certain kind of action. For instance, in order to do a good action, I have to presuppose that "good" is a value.

Page 11.

MESSERI: In the second paragraph of page 11, what do you mean by bias? How does bias distinguish itself from the value system?

MITTER: I don't know if I can give you a general description of that. One of the issues that I have not touched upon here is the issue of complexity. When we recognize in a natural scene we hardly ever make a mistake. Now, if we try to build a machine which recognizes objects in a natural scene, we are immediately faced with very complex algorithms.

What seems to happen when you work at a formal system level—when I say a formal system I mean some formal construction of models—you are faced immediately with exponential complexity. The only way that I can see of cutting down this complexity is that one may have the a priori knowledge, bias in some way, to reduce this complexity. Maybe you have sufficient initial knowledge which helps you in your task. Call that "bias," which helps you to explore the infinity of possibilities that you have, that you might be faced with in this algorithmic process and makes the process finite. But bias in the algorithm arises from a value system.

MESSERI: That is different from the way social science defines bias.

MITTER: So, bias here is used in a sense of prior knowledge.

ARONSON: You say, if I understand you, that to reduce the complexity of a system in some way, one must bias it. I would argue, based on the emphasis on time scales in your paper, that one way of reducing complexity is simply to explicitly establish the hierarchies of complexity within which the system concerned occurs. You may do that, if you wish, with some set of values, and certainly with some degree of subjectivity. But if you establish not only your time scale, but also your spatial scales, and, therefore, begin to have an explicit hierarchy of complexity, that simplifies the problem of analysing the system.

MITTER: The point is that when you talk about hierarchy, you mean presumably some kind of ordering. What do you mean by hierarchy?

ARONSON: In biology we speak of a system that occurs within a field, and that field occurs within a watershed, and the watershed occurs within a bio-geographical region, and between those three rungs, below the first and above the third, you could add many other rungs of biological, ecological and spatial complexity.

MITTER: I am saying that is one way of reducing complexity. You are quite right. The idea is to view a system at different levels of abstraction.

AUSENDA: You believe that "the optimality metaphor does not apply to the workings of an adaptive system...[because] it does not try...to optimize anything in particular...because of conflicting incommensurable criteria being present right in a single overarching criteria?" Do you consider the maximization of effectiveness as an optimizing value or not? If not, do you think that systems tend to optimize effectiveness?

MITTER: I don't have a very objective answer to this question. My sense says adaptive systems are not optimizing any criterion.

AUSENDA: So systems do not try to optimize effectiveness.

MITTER: No. Certainly from our knowledge of control we find that optimizing solutions are often very sensitive. You change the parameters of the ambient conditions and optimizing solutions tend to be highly non-robust against the changes. I think I said that it is a highly dynamic concept. One of the things that certainly I am trying to emphasize, is the dynamical view of systems. From what I know of theoretical economics, there a highly equilibrium-based and static view of a system is adopted. It seems to be totally inappropriate to do that if you want to talk about adaptiveness or robustness against some uncertainty or any of these issues.

ARONSON: As I suggested in my paper, we all inherit a balance of nature paradigm, in the Western world at least, and it is totally false. It does not correspond to biological reality. So, we all need to work on a changing paradigm, economists included. Ecologists speak today of a "flux of nature" paradigm. Extinction is the norm, not survival. It is a different idea from what we usually think. More species die out than those which survive. This is true of individuals as well, and probably most ecosystems. There may be selection among ecosystems, and we only see the survivors. Or we may see those that are in different stages of senescence and extinction. And others that are emerging, but they are certainly not in equilibrium, per se.

VOICE: What time frame do you use?

ARONSON: It depends on the question you are asking. If you study a population, then you have to set a time frame relevant to the human life span or the life span of the organism that you are concerned with at the population level. If, however, you are concerned with communities or ecosystems that last hundreds of generations, well then your time scale must be appropriate to that subject of study. Page 12.

THOMPSON: In your analysis of the coherence of classical music, you said it was effective because it was coherent.

MITTER: I would like to tell you that. There is a book by Schönberg on harmony. I think the title is "Harmony, Form and Coherence" (Schönberg 1994), and it is a very interesting book, where he talks about issues of coherence in music.

THOMPSON: Where do you think that "coherence" appplies to the individual or social construct that the individual comes from?

MITTER: I cannot answer that question. I don't know. What I can say is that, in limited contexts, one can postulate certain models and then test them out. I think what I am trying to get at here is that current paradigms of control and communication are quite inadequate to address many of the issues that seem to arise in the context of complex systems. And I am suggesting the notion of "coherence" as a paradigm, without giving any precise definition, as a principle for looking at examples. The use of "coherence" as a principle is in the same sense as the use of "fitness" as a principle in Darwinian evolution.

References

Arrow, K. T.
1974 See References at end of paper.
Changeux, J.-P., & A. Connes
1989 See References at end of paper.
De Finetti, B.
1974 *Theory of Probability*, Vols. 1 & 2. London/New York: John Wiley & Sons.
Habermas, J.
1973 *Erkentniss und Interesse*. [2nd. edition]. Frankfurt: Suhrkamp.
Schönberg, A.
1994 See References at end of paper.
Weyl , H.
1989 See References at end of paper.
Whitehead, A. N.
1978 See References at end of paper.