Proceedings of the American Philosophical Society, Vol. 106, No. 6. (Dec. 12, 1962), pp. 467-482.

THE ARCHITECTURE OF COMPLEXITY

HERBERT A. SIMON*

Professor of Administration, Carnegie Institute of Technology

A NUMBER of proposals have been advanced in recent years for the development of "general systems theory" which, abstracting from properties peculiar to physical, biological, or social systems, would be applicable to all of them. We might well feel that, while the goal is laudable, systems of such diverse kinds could hardly be expected to have any nontrivial properties in common. Metaphor and analogy can be helpful, or they can be misleading. All depends on whether the similarities the metaphor captures are significant or superficial.

I shall not undertake a formal definition of "complex systems." Roughly, by a complex system I mean one made up of a large number of parts that interact in a nonsimple way. In such systems, the whole is more than the sum of the parts, not in an ultimate, metaphysical sense, but in the important pragmatic sense that, given the properties of the parts and the laws of their interaction, it is not a trivial matter to infer the properties of the whole.

The four sections that follow discuss four aspects of complexity. The first offers some comments on the frequency with which complexity takes the form of hierarchy - the complex system being composed of subsystems that, in turn, have their own subsystems, and so on. The second section theorizes about the relation between the structure of a complex system and the time required for it to emerge through evolutionary processes: specifically, it argues that hierarchic systems will evolve far more quickly than non-hierarchic systems of comparable size. The third section explores the dynamic properties of hierarchically-organized systems, and shows how they can be decomposed into subsystems in order to analyze their behavior. The fourth section examines the relation between complex systems and their descriptions.

Thus, the central theme that runs through my remarks is that *complexity frequently takes the form of hierarchy, and that hierarchic systems have some common properties that are independent of their specific content*. Hierarchy, I shall argue, is one of the central structural schemes that the architect of complexity uses.

HIERARCHIC SYSTEMS

By a hierarchic system, or hierarchy, I mean a system that is composed of interrelated sub-systems, each of the latter being, in turn, hierarchic in structure until we reach some lowest level of elementary subsystem. *In most systems in nature, it is somewhat arbitrary as to where we leave off the partitioning, and what subsystems we take as elementary*. Physics makes much use of the concept of "elementary particle" although particles have a disconcerting tendency not to remain elementary very long. Only a couple of generations ago, the atoms themselves were elementary particles; today, to the nuclear physicist they are complex systems. For certain purposes of astronomy, whole stars, or even galaxies, can be regarded as elementary subsystems. In one kind of biological research, a cell may be treated as an elementary subsystem; in another, a protein molecule; in still another, an amino acid residue.

Just why a scientist has a right to treat as elementary a subsystem that is in fact exceedingly complex is one of the questions we shall take up. For the moment, we shall accept the fact that scientists do this all the time, and that if they are careful scientists they usually get away with it.

Etymologically, the word "hierarchy" has had a narrower meaning than I am giving it here. The term has generally been used to refer to a complex system in which each of the subsystems is subordinated by an authority relation to the system it belongs to. More exactly, in a hierarchic formal organization, each system consists of a "boss" and a set of subordinate subsystems. Each of the subsystems has a "boss" who is the immediate subordinate of the boss of the system.

We shall want to consider systems in which the relations among subsystems are more complex than in the formal organizational hierarchy just described. We shall want to include systems in which there is no relation of subordination among subsystems. (In fact, even in human organizations, the formal hierarchy exists only on paper; the real flesh-and-blood organization has many inter-part relations other than the lines of formal authority.)

SOCIAL SYSTEMS

I have already given an example of one kind of hierarchy that is frequently encountered in the social sciences: a formal organization. Business firms, governments, universities all have a clearly visible parts-within-parts structure. But formal organizations are not the only, or even the most common, kind of social hierarchy. Almost all societies have elementary units called families, which may be grouped into villages or tribes, and these into larger groupings, and so on. If we make a chart of social interactions, of who talks to whom, the clusters of dense interaction in the chart will identify a rather well-defined hierarchic structure. The groupings in this structure may be defined operationally by some measure of frequency of interaction in this sociometric matrix.

BIOLOGICAL AND PHYSICAL SYSTEMS

The hierarchical structure of biological systems is a familiar fact. Taking the cell as the building block, we find cells organized into tissues, tissues into organs, organs into systems. Moving downward from the cell, well-defined subsystems-for example, nucleus, cell membrane, microsomes, mitochondria, and so on-have been identified in animal cells.

The hierarchic structure of many physical systems is equally clear-cut. I have already mentioned the two main series. At the microscopic level we have elementary particles, atoms, molecules. macromolecules. At the macroscopic level we have satellite systems, planetary systems, galaxies. Matter is distributed throughout space in a strikingly non-uniform fashion. The most nearly random distributions we find, gases, are not random distributions of elementary particles but random distributions of complex systems, i.e. molecules.

A considerable range of structural types is subsumed under the term hierarchy as I have defined it. By this definition, a diamond is hierarchic, for it is a crystal structure of carbon atoms that can be further decomposed into protons, neutrons, and electrons. However, it is a very "flat" hierarchy. in which the number of first-order sub-systems belonging to the crystal can be indefinitely large. A volume of molecular gas is a flat hierarchy in the same sense. In ordinary usage, we end to reserve the word hierarchy for a system that is divided into a small or moderate number of subsystems, each of which may be further sub-divided. Hence, we do not ordinarily think of or refer to a diamond or a gas as a hierarchic structure. Similarly, a linear polymer is simply a chain, which may be very long, of identical sub-parts, the monomers. At the molecular level it is a very flat hierarchy.

In discussing formal organizations, the number of subordinates who report directly to a single boss is called his span of control. I will speak analogously of the span of a system, by which I shall mean the number of subsystems into which it is partitioned. Thus, a hierarchic system is flat at a given level if it has a wide span at that level. A diamond has a wide span at the crystal level, but not at the next level down, the molecular level.

There is one important difference between the physical and biological hierarchies, on the one hand, and social hierarchies, on the other. Most physical and biological hierarchies are described in spatial terms. We detect the organelles in a cell in the way we detect the raisins in a cake- they are "visibly" differentiated substructures localized spatially in the larger structure. On the other hand, we propose to identify social hierarchies not by observing who lives close to whom but by observing who interacts with whom. These two points of view can be reconciled by defining hierarchy in terms of intensity of interaction, but observing that *in most biological and physical systems relatively intense interaction implies relative spatial propinquity*. One of the interesting characteristics of nerve cells and telephone wires is that they permit very specific strong interactions at great distances. To the extent that interactions are channeled through specialized communications and transportation systems, spatial propinquity becomes less determinative of structure.

SYMBOLIC SYSTEMS

One very important class of systems has been omitted from my examples thus far: systems of human symbolic production. A book is a hierarchy in the sense in which I am using that term. It is generally divided into chapters, the chapters into sections, the sections into paragraphs, the paragraphs into sentences, the sentences into clauses and phrases, the clauses and phrases into words. We may take the words as our elementary units, or further subdivide them, as the linguist often does, into smaller units. If the book is narrative in character, it may divide into "episodes" instead of sections, but divisions there will be.

The hierarchic structure of music, based on such units as movements, parts, themes, phrases, is well known.

In social as in physical systems there are generally limits on the simultaneous interaction of large numbers of subsystems. In the social case, these limits are related to the fact that a human being is more nearly a serial than a parallel information-processing system. He can carry on only one conversation at a time, and although this does not limit the size of the audience to which a mass communication can be addressed, it does limit the number of people simultaneously involved in most other forms of social interaction. Apart from requirements of direct interaction, most roles impose tasks and responsibilities that are time consuming. One cannot, for example, enact the role of "friend" with large numbers of other people.

It is probably true that *in social as in physical systems, the higher frequency dynamics are associated with the subsystems, the lower frequency dynamics with the larger systems*. It is generally believed, for example, that the relevant planning horizon of executives is longer the higher their location in the organizational hierarchy. It is probably also true that both the average duration of an interaction between executives and the average interval between interactions is greater at higher than at lower levels.

SUMMARY: NEAR DECOMPOSABILITY

We have seen that hierarchies have the property of near-decomposability. Intra-component linkages are generally stronger than intercomponent linkages. This fact has the effect of separating the high-frequency dynamics of a hierarchy, involving the internal structure of the components, from the low frequency dynamics, involving interaction among components. We shall turn next to some important consequences of this separation for the description and comprehension of complex systems.

THE DESCRIPTION OF COMPLEXITY

If you ask a person to draw a complex object- e.g., a human face-he will almost always proceed in a hierarchic fashion. First he will outline the face. Then he will add or insert features: eyes, nose, mouth, ears, hair. If asked to elaborate, he will begin to develop details for each of the features - pupils, eyelids, lashes for the eyes, and so on - until he reaches the limits of his anatomical knowledge. His information about the object is arranged hierarchically in memory, like a topical outline.

When information is put in outline form, it is easy to include information about the relations among the major parts and information about the internal relations of parts in each of the suboutlines. Detailed information about the relations of subparts belonging to different parts has no place in the outline and is likely to be lost. The loss of such information and the preservation mainly of information about hierarchic order is a salient characteristic that distinguishes the drawings of a child or someone untrained in representation from the drawing of a trained artist.

NEAR DECOMPOSABILITY AND COMPREHENSIBILITY

From our discussion of the dynamic properties of nearly decomposable systems, we have seen that comparatively little information is lost by representing them as hierarchies. Subparts belonging to different parts only interact in an aggregative fashion - the detail of their interaction can be ignored. In studying the interaction of two large molecules, generally we do not need to consider in detail the interactions of nuclei of the atoms belonging to the one molecule with the nuclei of the atoms belonging to the other. In studying the interaction of two nations, we do not need to study in detail the interactions of each citizen of the first with each citizen of the second.

The fact, then, that many complex systems have a nearly decomposable, hierarchic structure is a major facilitating factor enabling us to understand, to describe, and even to "see" such systems and their parts. Or perhaps the proposition should be put the other way around. If there are important systems in the world that are complex without being hierarchic, they may to a considerable extent escape our observation and our understanding. Analysis of their behavior would involve such detailed knowledge and calculation of the interactions of their elementary parts that it would be beyond our capacities of memory or computation.

I shall not try to settle which is chicken and which is egg: whether we are able to understand the world because it is hierarchic, or whether it appears hierarchic because those aspects of it which are not elude our understanding and observation. I have already given some reasons for supposing that the former is at least half the truth - that evolving complexity would tend to be hierarchic - but it may not be the whole truth.

SIMPLE DESCRIPTIONS OF COMPLEX SYSTEMS

If a complex structure is completely unredundant-if no aspect of its structure can be inferred from any other-then it is its own simplest description. We can exhibit it, but we cannot describe it by a simpler structure. The hierarchic structures we have been discussing have a high degree of redundancy, hence can often be described in economical terms. The redundancy takes a number of forms, of which I shall mention three:

1. Hierarchic systems are usually composed of only a few different kinds of subsystems, in various combinations and arrangements. A familiar example is the proteins, their multitudinous variety arising from arrangements of only twenty different amino acids. Similarly, the ninety-odd elements provide all the kinds of building blocks needed for an infinite variety of molecules. Hence, we can construct our description from a restricted alphabet of elementary terms corresponding to the basic set of elementary subsystems from which the complex system is generated.
2. Hierarchic systems are, as we have seen, often nearly decomposable. Hence only aggregative properties of their parts enter into the description of the interactions of those parts. A generalization of the notion of near-decomposability might be called the "empty world hypothesis" -- most things are only weakly connected with most other

things; for a tolerable description of reality only a tiny fraction of all possible interactions needs to be taken into account. By adopting a descriptive language that allows the absence of something to go unmentioned, a nearly empty world can be described quite concisely. Mother Hubbard did not have to check off the list of possible contents to say that her cupboard was bare.

3. *By appropriate "recoding," the redundancy that is present but unobvious in the structure of a complex system can often be made patent.* The most common recoding of descriptions of dynamic systems consists in replacing a description of the time path with a description of a differential law that generates that path. The simplicity, that is, resides in a constant relation between the state of the system at any given time and the state of the system a short time later. Thus, the structure of the sequence, 1 3 5 7 9 11 ... , is most simply expressed by observing that each member is obtained by adding 2 to the previous one. But this is the sequence that Galileo found to describe the velocity at the end of successive time intervals of a ball rolling down an inclined plane.

It is a familiar proposition that the task of science is to make use of the world's redundancy to describe that world simply. I shall not pursue the general methodological point here, but shall instead take a closer look at two main types of description that seem to be available to us in seeking an understanding of complex systems. I shall call these state description and process description, respectively.

STATE DESCRIPTIONS AND PROCESS DESCRIPTIONS

"A circle is the locus of all points equidistant from a given point." "To construct a circle, rotate a compass with one arm fixed until the other arm has returned to its starting point." It is implicit in Euclid that if you carry out the process specified in the second sentence, you will produce an object that satisfies the definition of the first. The first sentence is a state description of a circle, the second a process description.

These two modes of apprehending structure are the warp and weft of our experience. Pictures, blueprints, most diagrams, chemical structural formulae are state descriptions. Recipes, differential equations, equations for chemical reactions are process descriptions. The former characterize the world as sensed; they provide the criteria for identifying objects, often by modeling the objects themselves. The latter characterize the world as acted upon; they provide the means for producing or generating objects having the desired characteristics.

The distinction between the world as sensed and the world as acted upon defines the basic condition for the survival of adaptive organisms. The organism must develop correlations between goals in the sensed world and actions in the world of process. When they are made conscious and verbalized, these correlations correspond to what we usually call means-end analysis. Given a desired state of affairs and an existing state of affairs, the task of an adaptive organism is to find the difference between these two states, and then to find the correlating process that will erase the difference.

Thus, problem solving requires continual translation between the state and process descriptions of the same complex reality. Plato, in the Meno, argued that all learning is remembering. He could not otherwise explain how we can discover or recognize the answer to a problem unless we already know the answer. Our dual relation to the world is the source and solution of the paradox. We pose a problem by giving the state description of the solution. The task is to discover a sequence of processes that will produce the goal state from an initial state. Translation from the process description to the state description enables us to recognize when we have succeeded. The solution is genuinely new to us- and we do not need Plato's theory of remembering to explain how we recognize it.

There is now a growing body of evidence that the activity called human problem solving is basically a form of means-end analysis that aims at discovering a process description of the path that leads to a desired goal. The general paradigm is: given a blueprint, to find the corresponding recipe. Much of the activity of science is an application of that paradigm: given the description of some natural phenomena, to find the differential equations for processes that will produce the phenomena.

SUMMARY: THE DESCRIPTION OF COMPLEXITY

How complex or simple a structure is depends critically upon the way in which we describe it. Most of the complex structures found in the world are enormously redundant, and we can use this redundancy to simplify their description. But to use it, to achieve the simplification, we must find the right representation.

The notion of substituting a process description for a state description of nature has played a cen-tral role in the development of modern science. Dynamic laws, expressed in the form of systems of differential or difference equations, have in a large number of cases provided the clue for the simple description of the complex.

CONCLUSION

Our speculations have carried us over a rather alarming array of topics, but that is the price we must pay if we wish to seek properties common to many sorts of complex systems. My thesis has been that one path to the construction of a non-trivial theory of complex systems is by way of a theory of hierarchy. Empirically, a large proportion of the complex systems we observe in nature exhibit hierarchic structure. On theoretical grounds we could expect complex systems to be hierarchies in a world in which complexity had to evolve from simplicity. In their dynamics, hierarchies have a property, near-decomposability, that greatly simplifies their behavior. Near-decomposability also simplifies the description of a complex system, and makes it easier to understand how the information needed for the development or reproduction of the system can be stored in reasonable compass.