

How Complex Are Complex Systems?

Herbert A. Simon

Carnegie-Mellon University

As it has turned out--and not, I think, by conscious design--this symposium has a truly Hegelian structure. The thesis and antithesis are provided by the themes of complexity and simplicity, respectively. Professor Nelson pleads the case for complexity, and I for simplicity; and we both count on Professor Suppes to propose the synthesis.

Methods of characterizing the simplicity or complexity of systems are many. Let me mention some of them:

1. Systems that have many components may be considered complex relative to systems that have few. Hence the *cardinality* of a set may be taken as one measure of its complexity.
2. Systems in which there is much *interdependence* among the components are generally regarded as more complex than systems with less interdependence among components.
3. Systems that are *undecidable* may be regarded as complex in comparison with those that are decidable.
4. The complexity of systems may be measured by their *information content*, in the Shannon-Wiener sense. By this criterion, systems with many identical components are less complex than systems of comparable size whose components are all different.

We can speak not only generally of the complexity of systems, but specifically of the complexity of theories or of problem domains:

5. Closely related to the notion of informational complexity is the idea of measuring the complexity of theories by the number of their *parameters*, or by the number of *symbols* required to state them.
6. There is today a very active interest among mathematicians and computer scientists in *computational complexity*--in the maximum

number or the expected number of elementary computational steps required to solve the problems in a given class.

7. Measures of *problem difficulty* may be regarded, at least under some circumstances, as a particular class of measures of computational complexity.¹

Complexity may lie in the structure of a system, but it may also lie in the eye of a beholder of that system. Even when a system is "inherently" simple--i.e., describable, in principle, in simple terms--an observer may fail to discover that simple description, and may be able to characterize the system only in a very complex way. Moreover, the simplicity or complexity of a description will depend upon the units that are selected as primitives. The description of a computer program in a high-level language like ALGOL will generally be much briefer than the description of the same program in machine language. The relativity of complexity to the set of primitives is especially evident in the cases of informational and computational complexity, but it also applies to cardinality--which depends on what is taken as a single element.

Finally, an exact description of a system will generally be more complex than a description of that system to some specified degree of approximation. We may regard the inherent complexity of a structure as defining a lower bound for the complexities of exact descriptions, but approximate descriptions may be even less complex than this lower bound.

This paper will be concerned with complexity mainly in the second of the senses listed above--the degree of interdependence of components. It will have something to say, too, about informational complexity, but relatively little about decidability, computational complexity, or the other kinds of complexity that have been mentioned.

1. Fortune Smiles Upon Kepler and Newton

One of the functions of the experimental method is to substitute simple artificial systems for the complex systems that Nature presents to us. Sometimes, however, as in celestial mechanics before Sputnik, it is not possible to perform experiments. Then we have no choice but to deal with complexity as we find it.

The solar system known to Kepler and Newton had a substantial number of components: a sun, six planets (counting the earth), and about ten visible satellites belonging to three of those planets--to say nothing of the comets darting in and out of the system and the background of fixed stars. If we consider only pairs of bodies, we would have to take into consideration $17 \times 16 = 272$ potential interactions among the elements of this system. There is no reason in principle, of course, why the behavior of the system should depend just upon interactions of pairs of elements, but for the moment, we will make the simplifying assumption that it does. The prospect of handling a coefficient matrix of 272 elements provides complexity enough. Kepler and Newton did no such thing. Kepler detected three observational regularities: the orbits of the planets about the sun

are ellipses with the sun at a focus, equal areas are swept out in a given planetary orbit in equal times, and the periodic times of the planets vary with the $3/2$ power of their distances from the sun. Newton showed that these phenomena, together with the analogous ones for the satellites, could be deduced from his Laws of Motion, taken in combination with the gravitational law, gravitational force varying inversely with the square of the distance between a pair of bodies. But the observed regularities, and the derivation as well, depend upon the assumptions that each particular orbit under consideration was determined by the interaction between a single central body (the sun or a planet) and a single orbiting body (a planet or satellite). In each case, interactions with all other components of the system were ignored, and yet an excellent fit was obtained between the theoretical derivations and the observations.

The fact that the simple calculations "worked" was due to the actual solar system being a very special case, if we regard it as drawn by sampling from some population of possible solar systems. If the deck of cards that Nature dealt to Kepler and Newton was not stacked, it was at least a very lucky deal.

First there was a single body, the sun, that was larger by three orders of magnitude than any other body in the system. Second, there were six bodies, the planets, that were several orders of magnitude larger than their satellites. Third, the distances of the planets from each other were of the same order of magnitude as their distances from the sun, while the distances of the satellites from their planets were orders of magnitude smaller than their distances from the sun. None of these distributional facts follows from the laws of mechanics, but had they not been true of the solar system, Kepler's regularities and Newton's derivations would not have described that system.² Although Newton's Laws are generally valid (up to the classical, non-relativistic approximation) for systems of masses, the relatively simple calculations used to test those laws would not have sufficed had the system been more "general." Because of these relative sizes and distances, each planet orbited around the sun almost exactly as if it had been attracted by the sun alone, and each satellite around its planet in a similar way.

This simplification depended not only upon the distribution of masses and distances, but upon the forms of the laws themselves. Because of the inverse relation between gravitational attraction and distance, interactions between pairs of bodies close to each other would be much larger, other things being equal, than interactions of bodies far apart. Each factor of three in distance would produce an order of magnitude difference in strength of attraction. Second, because the attractions between a body and two others depended only on the directions and distances of the latter two from the former, the force per unit of mass exerted by the sun on each planet was almost exactly equal to the force exerted upon each of that planet's satellites. Since, in addition, accelerations were additive by the parallelogram law, the motion of a satellite *relative* to its planet could be calculated without considering the action of the sun on either of them.

For all of these reasons, the solar system turned out to be far simpler than one might conjecture simply by noting that it involves 272 pairwise interactions. Without denigrating the accomplishments of Copernicus, Kepler, and Newton, we may say that its scrutability hinged on its being less complex than it might have been, had Nature or Fortune, whichever we are to hold responsible, been less kind.

2. Some Causes for Simplicity

The solar system is not by any means the only example of a natural system that is simpler than it might be--in which only a small fraction of the potential interactions between pairs of elements are of any significance. In the solar system, we did not find any convincing reason why it should be as simple as it is, although the rapid attenuation of force with distance played an important role in limiting the number of strong interactions. In a number of the other cases we will examine, we will find that there are sufficient causes for the simplicity.

It will be convenient to represent the interconnectedness of a system by an incidence matrix, a matrix of zeros and ones, the $(i,j)^{th}$ element being 1 if the i^{th} element interacts with the j^{th} , and 0 otherwise. There are a number of different reasons why we might expect most real-world systems to have rather sparse incidence matrices. Some of these reasons are quite general, others relate to specific classes of systems. Let us begin with the general and move toward the specific.

1. It's an Illusion. It is true that most of the complex systems about which much is known have rather sparse incidence matrices, but this may simply be because systems lacking this property are inscrutable and there is little we can say about them. There is undoubtedly some merit in this claim. We attend to things about which there is something to say--that is, that are simple enough to yield to analysis. But this is not the whole story.

2. Spatial Relations. In most physical and biological systems, as was the case of the solar system, the strengths of interactions decrease with distance. Each element can have strong interactions only with the relatively few elements that are close to it. Consider, for example, a system that is essentially one-dimensional, so that its elements can be arranged along a line. Suppose that the behavior of each element can be described by a differential equation, whose independent variables are the positions (and perhaps other characteristics) of the other elements. If we now display the matrix of coefficients of these equations, with the rows and columns ordered as the elements are, the large entries will be close to the diagonal, for these near-diagonal entries represent the interactions among elements that are close neighbors. The differential equations and partial differential equations of physical systems typically have this near-diagonal structure. If the system is also homogeneous, so that successive rows measured from the diagonal are essentially the same, then the system is usually amenable to one form or another of analysis. Such analysis is the stock in trade of books on methods of mathematical physics.

3. Evolution Prefers Hierarchy. For present purposes, we may think of a hierarchy as a system in which the elements are arranged in a tree, so that each element is influenced by a single ancestral element, and influences, in turn, a number of descendant elements. Later, we will subsume a somewhat broader class of systems under this term. The behavior of hierarchies, we shall see, is much simpler than the behavior of systems with comparable numbers of elements in which each element interacts with every other element. We shall also see that evolutionary processes are much more likely to produce hierarchies than non-hierarchical systems of comparable size.

4. Information Demands. When forces, like gravitational forces, impinge on the elements of a system, there may be no in-principle limit upon their number (other than the limits, discussed earlier, of distance when force depends on propinquity). Where, however, interaction involves some kind of expenditure of energy or attention, either by the element that exerts an influence, or by the one that is subject to it, then this scarce resource may have to be allocated, with a resulting limit on the amount of interaction that is possible. This is a major factor limiting the complexity of interaction in biological and social systems, especially when the elements must process outgoing or incoming information serially, one or a few items at a time.

5. Comprehensibility. The amount of interaction in a system may be limited not only by comprehensibility to the viewer (the first point above), but by comprehensibility to the system itself. For example, in the case of a system that learns--that is, changes its own structure adaptively--some understanding of that structure may be required in order to identify the substructures in which adaptive changes should be made, and to characterize the nature of the changes.

There is an impressive number of reasons, therefore, why we should expect most things in the world not to interact strongly with most other things. Moreover, in most cases, we would expect the incidence matrix describing the interactions among elements not only to be sparse, but also to have some definite structure. We have seen that one form this structure may take is to concentrate the larger interactions close to the diagonal of the matrix. Another form is to arrange the interactions hierarchically. The theory of "possible" or "plausible" forms of complex systems is still in a primitive state, but already these two classes of systems emerge as important special cases.

3. Nearly-Decomposable Systems

Elsewhere,³ I have discussed at some length systems whose incidence matrices are approximately hierarchic in character, giving both some reasons why this configuration is found so frequently in natural systems, and some consequences of a system's having this form. I will summarize the arguments briefly here.

First, I must specify more carefully the class of systems I wish to consider. Suppose that the strengths of interactions among pairs of elements in the system vary (as they do in the solar system) over several

orders of magnitude, clustering, say, into "strong," "moderate," and "weak" interactions, with large gaps between the three classes. Suppose, further, that we can order the elements in such a way that all the strong interactions are arranged in non-intersecting blocks along the diagonal of the incidence matrix. Suppose that these blocks are enclosed in larger along-diagonal blocks that contain all of the moderate interactions, so that only weak interactions appear in the remaining portions of the matrix. Under these conditions, we will say that the system is *nearly completely decomposable*. (See Figure 1.)

	1	2	3	4	5	6	7	8
1	x	S	W	M	-	-	W	W
2	S	x	S	W	M	-	W	-
3	M	S	x	W	M	-	-	W
4	M	M	-	x	S	W	W	W
5	W	M	M	S	x	-	W	W
6	W	-	-	W	W	x	S	S
7	-	W	W	W	W	S	x	S
8	-	W	-	W	W	M	S	x

S = Strong; M = Moderate; W = Weak

Figure 1

A Nearly-Completely-Decomposable System

Notice that strong interactions occur only within the three sets of elements (1 2 3), (4 5), and (6 7 8). Moderate interactions also occur between elements of the first two of these sets. All the remaining interactions are weak. We can summarize the structure of the matrix by depicting the elements as nesting sets: (((1 2 3)(4 5))((6 7 8))), where numbers of parentheses indicate the degree of "insulation" of each set of elements from adjacent sets. The parenthesis notation shows clearly how the sets can be arranged hierarchically. The entire group of eight elements is divided into two sets (elements 1 through 5, and 6 through 8, respectively), and the first of these sets is again subdivided (elements 1 through 3 and 4 through 5, respectively).

Now consider the assembly of systems from stable subsystems that already exist. Suppose that new systems can be formed by bringing two or more existing systems into propinquity with each other. Suppose further that only a small fraction of such potential combinations form stable assemblies, the remainder dispersing almost as soon as they are formed. The rate at which large, complex systems will evolve then becomes a matter of how often components capable of forming a stable combination will be found together in a single neighborhood. Clearly, the probability of such an event will fall off rapidly with the number of components simultaneously involved. Hence, a new system of a given size is more likely to be formed from a

small number of medium-size systems than directly from a large number of small systems. As the process of composition proceeds, the larger systems that emerge will exhibit more and more layers of substructure.

Under the simplest assumptions, the expected time for a large system to assemble spontaneously from its "atomic" components increases geometrically with the number of those components. On the other hand, if there is a hierarchy of potential "subassemblies," with about the same number of components at each level, then the expected time required to assemble a system of n elements will grow only with $\log n$. That is to say, when large systems are formed in this hierarchic way, the doubling time for the size of the largest systems that evolve will be roughly constant. Systems of size 100 might be expected to evolve in an environment of systems of size 50 in about the same time as systems of size 50 in an environment of systems of size 25.

The argument just outlined contains a number of "iffy" assumptions. Perhaps the most crucial is the assumption that the probability that a set of components brought into propinquity with each other will form a stable system is independent of the size of the components. But even this assumption could be considerably modified without destroying the evolutionary advantage of hierarchic over non-hierarchic structures. The argument provides a plausible explanation for the observed fact that hierarchy is almost universal among natural structures--for example, the hierarchy that extends from multi-celled organisms through cells, molecules, atoms, and nuclei to the "elementary" subnuclear particles which themselves are more and more revealing hierarchic structure.

It is a common, but not logically necessary, characteristic of hierarchic systems that they are nearly-decomposable, in the sense in which that term was defined earlier. That is to say, as we proceed upward, from level to level, in the hierarchy, the strengths of the interactions between elements belonging to different components become weaker and weaker. For example, the energies of typical atomic covalent bonds are in the range of 80-100 k calories per mole, ionic bonds slightly lower, hydrogen bonds₄ in organic molecules, 10 k cal/mole, and van der Waals bonds still lower. Weaker bonds, in turn, are generally associated with slower, lower-frequency, interactions.

As a matter of fact, if the hierarchy emerges through an evolutionary process like that described above, we would expect some such relation of the bond strengths at different hierarchic levels, for this arrangement would allow the weaker, high level bonds to form without disturbing the stronger bonds of the stable subsystems of which the larger systems were composed.

Let us suppose that we have a hierarchic system whose dynamic behavior can be represented by a system of first-order linear differential equations with constant coefficients. This is, of course, a very special case, but it will illustrate some very fundamental qualitative characteristics of the behavior of hierarchic systems. The matrix of Figure 1 can now be taken to represent the matrix of coefficients of these differential equations. If we solve the equations--say, for a system of n elements--we will obtain

solutions of the form:

$$(1) \quad x_i = a_{ij} \exp(\lambda_j t),$$

where the summation runs from 1 to n , and the a_{ij} and λ_j are constants.

The λ s are called the *characteristic roots* of the equations; if their real parts are negative, the corresponding terms will go to zero as t increases, if positive, the terms will increase without limit. Hence, stability of the dynamic system requires that the real parts of all the characteristic roots be negative.

If the matrix is nearly decomposable, as in Figure 1, then approximate solutions can be found for (1) without solving the entire system of equations simultaneously, but by proceeding, in stages, to solve some much smaller systems. First, we replace all the moderate and weak interactions by zeroes, so that the resulting systems is decomposed into independent subsystems, viz:

	1	2	3	4	5	6	7	8
1	x	S	-	-	-	-	-	-
2	S	x	S	-	-	-	-	-
3	-	S	x	-	-	-	-	-
4	-	-	-	x	S	-	-	-
5	-	-	-	S	x	-	-	-
6	-	-	-	-	-	x	S	S
7	-	-	-	-	-	S	x	S
8	-	-	-	-	-	-	S	x

Now we can solve the systems consisting of (1 2 3), (4 5), and (6 7 8) independently of each other. It can be shown⁵ that the behavior of these independent subsystems corresponds to the short-run, high-frequency behavior of the entire system. Further, for purposes of studying the longer-term lower-frequency behavior of the larger system, we can compress each of the three submatrices into a single scalar variable (an aggregate index of the behavior of the subsystem), and analyse the behavior of the resulting aggregated system:

	A	B	C
A	x	M	W
B	M	x	W
C	W	W	x

where A corresponds to the subsystem (1 2 3), B to the subsystem (4 5), and C to the subsystem (6 7 8). Aggregates of the moderate and weak linkages between these subsystems have been reintroduced into the aggregated matrix. We can now repeat the same trick again, ignoring the weak

interactions to analyze separately the two components C and $D = (A \ B) = ((1 \ 2 \ 3) \ (4 \ 5))$. Solution of the systems D and C will now provide a description of the middle-run dynamics of the system. Finally, reintroducing the weak interactions, we obtain the system:

$$\begin{array}{cc} & D \quad C \\ D & x \quad W \\ C & W \quad x \end{array}$$

The behavior of this aggregated system describes the long-term, low-frequency behavior of the original disaggregated system.

The near-decomposability of the physical universe is the principal basis for its comprehensibility, and we can now see the fundamental reason for this fact. In any nearly-decomposable system, we can select a particular frequency band and study the behavior of the system just at the level of aggregation corresponding to that frequency range. Variables moving only at lower frequencies can be treated as constants; and the weak interactions of a system with other systems, which correspond to these lower frequencies, can be ignored. Variables interacting within subsystems at higher frequencies can be aggregated, because their short-run dynamics will settle them down to equilibrium during the intervals of interest to us. Hence, if our world is complex, it is not nearly as complex as it might be. It is so arranged that we do not have to deal with all of its interaction at once.

A further source of simplification of the kinds of hierarchic systems that are found in nature or produced by art is that the components at any given level often belong to a relatively small number of species, the elements of a given species being either identical or at least similar in important respects. The set of types of components that are used as the building blocks at any level in a hierarchic system may be regarded as the *alphabet* for that level. Thus, atoms are built from the alphabet of a few dozen (perhaps) elementary particles; molecules from the alphabet of a hundred kinds of atoms; DNA from the alphabet of four nucleotides; proteins from the alphabet of 20 amino acids; and so on.

Limiting the building blocks of a system to a small alphabet produces simplicity by redundancy. To understand the microstructure, it is not necessary to analyse each substructure, but only representative members of each *type* of substructure. Thus, we construct a theory of the carbon atom, and then use that theory in calculating the behavior of the myriad of organic systems in which chains of carbon atoms constitute the backbone.

Even when the building blocks are not identical, we may be able to construct a general theory that accounts for much of their behavior. One of the most important examples is cell theory, or the somewhat more specific theories of particular classes of cells, like neurons. Other important examples of theories that simplify when we postulate a typology of components can be found in the social sciences: notably economic theories that build on such units as "economic man" or "the representative firm."

Some interesting results have been obtained that show that the dynamic stability of a wide class of complex systems depends on the connexity of their parts--for low levels of connexity they will be stable, but above a certain critical level they will suddenly become unstable. This phenomenon was demonstrated by simulation by Gardner and Ashby, and subsequently proved analytically by May.

Following May, we consider a system with n variables that obeys some nonlinear set of first order differential equations, and we characterize its behavior in the neighborhood of an equilibrium point by the linear terms of a Taylor's expansion:

$$dx/dt = Ax,$$

where x is a vector, and A a matrix of coefficients. The system described by the equations will be stable at that equilibrium point if and only if all the eigenvalues of A have negative real parts. Now suppose that the coefficients a_{ij} are chosen randomly in such a way that, with probability $(1-C)$, they are zero, and with probability C , non-zero. Suppose, further, that the non-zero elements have a distribution whose mean value is zero and mean square value α . Then it can be shown that the system will be almost certainly stable (with probability approaching one) if

$$\alpha(nC)^{1/2} < 1$$

and almost certainly unstable otherwise.

From this result it appears that, among highly interconnected systems, only very special systems will be stable. Since unstable systems will have only a transient existence, we may expect most systems that are actually observed to have relatively low levels of connexity of their parts.

4. Other Bases for Aggregation

In a great many complex systems, analysis is much simplified by the fact that the influences upon each element of all the other elements *aggregate* in a simple way. The simplest such law of aggregation is numerical addition. Thus, for example, the total weight on a scale is simply the sum of the weights of the objects placed on it. The net effect does not depend at all on the numbers of such objects or the distribution of their individual weights.

A slightly more complicated law of the aggregation of influences is the parallelogram of forces. Here it is the components along each dimension that are independently additive. In our earlier example of the solar system, we saw how simple laws for the aggregation of forces permitted, for example, spherical bodies to be replaced in the analysis by point masses.

Through simple laws for the aggregation of influences, the detailed sources of influence upon each element lose their identity. *Market systems*

exhibit a remarkable simplicity based on this loss of identity of specific interactions. Suppose that we have a market with n suppliers and m buyers, where n and m may be very large numbers. Suppliers wish to maximize their profits, and buyers their utilities. Each supplier can determine the optimal amount of the commodity in question to deliver to the market on the basis of information about one parameter--the price. And each buyer can determine the amount he should buy on the basis of the same parameter. If the market incorporates an appropriate tatonnement mechanism, then an excess of supply over demand will lower price, influencing buyers to buy more and sellers to supply less. Similarly, an excess of demand over supply will increase price, influencing buyers to buy less and sellers to supply more. If the adjustment reactions are not too volatile, the price will settle to an equilibrium that will just clear the market.

The special point of interest in this familiar Invisible Hand is its economy of information. The offers, S_i , of individual suppliers are all pooled, by addition, into an aggregate quantity, the total supply, Q_s . The quantities demanded by buyers, B_i , are aggregated in the same way into Q_d . A single parameter, the price, P_r , changes under the influence of the difference between these two aggregates. The behavior of each supplier and buyer, and hence the individual supply and demand quantities, change under the influence of this single parameter. For a market of n sellers and m buyers, we have this kind of incidence matrix for the system:

	S1	S2	..	S _n	Q _s	P _r	Q _d	B1	B2	..	B _m
S1	x	-	..	-	-	x	-	-	-	..	-
S2	-	x	..	-	-	x	-	-	-	..	-
..	..	..	..	..	..	..	..	..	..	..	..
S _n	-	-	..	x	-	x	-	-	-	..	-
Q _s	x	x	..	x	x	-	-	-	-	..	-
P _r	-	-	..	-	x	x	x	-	-	..	-
Q _d	-	-	..	-	-	x	x	x	x	..	x
B1	-	-	..	-	-	x	-	x	-	..	-
B2	-	-	..	-	-	x	-	-	x	..	-
..	..	..	..	..	..	..	..	..	..	..	..
B _m	-	-	..	-	-	x	-	-	-	..	x

where the i^{th} variable is influenced by the j^{th} if the ij^{th} entry in the matrix is x .

Notice that this sort of influence mechanism leads not only to a very sparse matrix of interactions, but also permits the behavior of the component subsystems to be described in simplified terms. Instead of having to write the supply function as $S_i(S_1, S_2, \dots, S_n, B_1, B_2, \dots, B_m)$, we can write simply, $S_i(P_r)$, with a corresponding simplification of the demand function to $D_i(P_r)$. To this we have to add an adjustment process of the general form, $dP_r = F(Q_d - Q_s)$. But now the total quantities supplied and

demanded will also be simply functions of P_r , $Q_s(P_r)$, and $Q_d(P_r)$, so that we can limit our analysis, if we wish, to the 3×3 submatrix representing the interactions of these three variables. Hence, the anonymity of the specific interactions leads to a tremendous simplification, not merely of the aggregate system, but of the functions describing the component systems as well.

Another major class of aggregation procedures rest on the *Law of Large Numbers*. As a major example of the simplifying potential of that law we have statistical mechanics, which turns the enormous complexity of a mechanical microsystem into the relative simplicity of macroscopic thermodynamic regularities.

With stochastic processes, we simplify the description of the history of a dynamic system by ignoring the time paths of its various elements and considering only their distribution among states of the system and their probabilities of transfer from one state to another. A further large step of simplification is achieved when we consider only the equilibria or steady states of such systems, which can often be derived with little or no concern for the paths whereby they reach those states.

I hardly need belabor the almost innumerable applications of the Law of Large Numbers in its various guises as a device for the simple, aggregative treatment of complex systems.⁶

5. The Complexity of Program Structures

Large computer programs are themselves complex structures. Skilled programmers have always understood that certain principles must be followed in constructing such programs if confusion is to be avoided and the resulting programs are to be debuggable. In the past few years some of these principles have been made explicit under the rubric of "structured programming."

The key to comprehensibility in large programs is that there not be excessive cross-references from one part of the program to the other parts--that is to say, that the interconnexity of the several parts of the program be kept within bounds. A principal technique for doing this is (1) to structure the program as a hierarchy of subroutines, each routine, in its execution, calling upon routines at lower levels of the hierarchy; and (2) to avoid, or at least strictly limit, references from one routine to another except through their explicit inputs and outputs. (The term "hierarchy" should not be taken too literally here, for some of the routines may be recursive, calling themselves.)

Each routine in a well-structured program can be characterized in terms of the inputs it requires, the outputs it produces, and the transformation it performs on the inputs. The internal code of the routine, the manner in which it carries out the transformation, is irrelevant to the characterization. Hence, any other routine in the system may call upon a given routine without having knowledge of its internal structure; all communication between routines is channeled through their inputs and outputs.

By the same token, any routine can be altered, provided that its inputs, outputs and transformation are not changed, without any effect upon the performance of the entire program. This modularization permits the program to be modified, say, to make it more efficient in execution, without each change propagating through the entire system and requiring additional changes to be made elsewhere. If we examine the dynamic behavior of such a system, we find very high rates of interaction among the parts of sub-routines at low levels of the hierarchy, and progressively lower rates of interaction among the subroutines at higher levels. The system has precisely the properties of a nearly-decomposable system that we discussed earlier.

It is instructive to see how such a program conserves the transmission of information by making use of it "locally" within particular subroutines, and then destroying all but the final output when the routine is completed. Suppose, for example, that we wrote a closed subroutine to perform 4×4 multiplications, using the same algorithm that we would use with paper and pencil. Then, the routine would produce successively the four products of the multiplicand by the individual digits of the multiplier, would hold on temporarily to these intermediate products, would then add them together to get the final answer, and finally would erase them. The process would take the 8 input digits. It would then produce an intermediate product of 4 or 5 digits, then another, bringing the total to 8 or 10, then another, and another. Storage would have to be provided for 16 to 20 intermediate digits just before the final addition was carried out, but these could be deleted afterwards, and only the final product of 8 or 9 digits retained.

The decisions of how to segment the program into subroutines are largely determined by these considerations of transmitting only a small amount of information between routines. A typical subroutine takes a small amount of input information, expands it into a larger amount of information required while carrying out the transformation, and then compresses it again into a small amount of output information. It is precisely at these junction points where the information that has been generated can be summed up in a few output symbols that it is profitable to segment the routines.

Programs are created by the conscious problem-solving activity of a programmer rather than by a process of biological evolution. (I speak of the present state of the art, and leave out of account programs that learn and automatic programming systems.) In such deliberately created system, the programmer's requirement for conceptual simplicity and clarity enforces hierarchy just as it is enforced by processes of natural selection in evolving systems. In both cases, the resulting structures will generally be hierarchic nearly-decomposable structures.

6. Conclusion

It has been the general theme of this paper that we have no cause to despair in the face of complexity, or to imagine a degree of complexity in nature that does not exist. We overestimate the complexity of real systems when we simply count the number of their elements and the number of possible interactions among those elements. The world is mostly empty or

sparsely populated; most of the potential interconnections of things are simply absent, or, if they exist, are negligible in importance. The systems that occur in nature are mostly hierarchic and nearly decomposable. Even when large numbers of elements exercise influences on other elements, those influences can most often be approximated by aggregating them in simple ways--by summing them arithmetically, for example.⁷

Simplicity, of course, is a matter of degree. The atmospheric system surrounding our Earth is sufficiently complex that, modeling it with our largest computers, we are still able to predict only the grossest features of our weather, and barely in real time, at that. The fastest computers, and the powerful tools of operations research and numerical analysis do not suffice to tell us all we want to know about some of the complex systems that are of concern to us.

Not all of our difficulties in understanding the world around us, however, can be attributed to complexity. Some of it stems not from the complexity of systems, but from the inscrutability of their elementary mechanisms. The recent disappointments we have had in the forecasting models used to guide economic policy *may* be due to the intrinsic complexity of the economic system. On the other hand, it may equally well be due to the fact that we incorporated the wrong assumptions in the simple models that we use.

The human species has survived and thrived in the world, simple or complex as it may be, not so much through the speed and power of its computational capacities, as by exploiting the fact that the systems of interest to it represent highly special cases that can often be analyzed by relatively simple means, provided their underlying structure is detected. This argues for a strategy of searching for that structure, of pattern induction--a skill that is rather highly developed in the animal kingdom--followed by special analysis and heuristic problem-solving search, rather than brute-force analysis of very general classes of highly interconnected complex systems.

Notes

¹In his comments, Professor Suppes reminds us that I have not exhausted the list of possible meanings for the term "complexity." He mentions specifically geometric complexity (the complexity that arises as we move from discrete to continuous systems), complexity as measured by logical depth in the sense of Hintikka, and complexity of proofs measured by structural characteristics of their graphs.

²There may, of course, have been additional lawful reasons why the solar system had to be just so--although even today we do not know what those reasons are.

³"The Architecture of Complexity," Chapter 4 in The Sciences of the Artificial. Cambridge, Mass.: M.I.T. Press, 1969.

⁴The Sciences of the Artificial, pp. 103-104. See also Diane M. Ramsey (ed.). Information and Control Processes in Living Systems. New York: New York Academy of Sciences (1967): 120-121.

⁵H. A. Simon and A. Ando, "Aggregation of Variables in Dynamic Systems," Econometrica 29(1961): 111-138.

⁶In his comments, Professor Suppes inquired about my views on the relation of complexity to randomness. The three preceding paragraphs constitute part of my answer. In recent years, Chaitin and others have developed a theory of random sequences that equates randomness with the impossibility of representing a sequence by a formula that is shorter than the sequence itself. Although, with this definition, the randomness of any particular sequence depends upon the specific coding scheme that is employed, interesting theorems have been obtained, holding for all coding schemes, regarding the relative frequency of random sequences among all possible sequences. Within this theory of random sequences, it can be seen that randomness means complexity.

⁷My only serious quarrel--but it *is* a serious quarrel--with Professor Nelson's excellent paper concerns the failure, as it seems to me, of the Lake Erie ecological model to take sufficient advantage of the various opportunities for simplification that are outlined here. My own experience with environmental modeling, gained as a member of the President's Science Advisory Committee and as chairman of the Committee on Air Quality Control of the National Academy of Sciences, persuades me that such matters are best analysed, for public policy purposes, with simple, crude models. I have at least five objections to the Lake Erie model. First, there seems to me to be no practical reason why the four principal submodels need to be closely linked; an aggregate model of the linkage should be fully adequate for the purposes of the inquiry. Second, the model actually constructed makes more subtle demands on our understanding of the underlying mechanisms (economic, chemical, or biological) than current scientific knowledge can probably support. Third, there is no practical need to run out year-by-year scenarios; steady-state analysis would seem to be quite enough and would be orders of magnitude simpler in terms of computation. Fourth, the policy instruments available for ameliorating the situation are sufficiently crude (if powerful) as to make it rather futile to model system details that cannot be modified independently and in corresponding detail. Fifth, ameliorative action can always be taken on an incremental basis with feedback, so that we need not predict with exactitude the magnitudes of the effects. A policy of cleaning up Lake Erie by identifying and eliminating major sources of phosphate, roughly ordered by cost of suppression and continuing until conditions in the Lake begin to improve rapidly, would not be far wide of the mark. Such a strategy has been highly successful in solving the eutrophication problems of Lake Menona and Lake Washington. Although engineering draws upon science, we should not confuse the two. If we wish to investigate scientifically the theory of algal bloom, we should study it in highly simplified and isolated subsystems. If we want to improve the environment, we should use highly aggregative, crude, simple models that are not more refined than the policy instruments we intend to apply. The whole thesis of my paper is that we do not need to study

everything at once in its full, intricate relation with everything else.
I would apply that thesis here to achieve a radical simplification of the
Lake Erie Model.