p. 12, p. 17 [start on page 6]

September 1, 1969

Typos: 3, 8, 12, 22, 26, 27, 28,

DESIGNING ORGANIZATIONS FOR AN INFORMATION-RICH WORLD

Herbert A. Simon
Carnegie-Mellon University

If men do not pour new wine into old bottles, they do something almost as bad--they invest old words with new meanings. "Work" and "energy" are venerable members of the English vocabulary words, but since the Industrial and Scientific Revolutions, they have acquired entirely new meanings, in scientific and technical contexts. They have become more abstract, divorcing themselves from directly sensed qualities of human activity; and they have become more precise, expressing themselves in specific, quantitative units of measurement--the foot-pound, the erg--and fundamental exact scientific laws--the Conservation of Energy.

Hence, the word "energy," uttered in a contemporary setting, may represent quite different concepts stored in memory and quite different processes of thought from the word "energy" uttered in the 18th century. That Burke's or Jefferson's vocabularies are similar to ours should not deceive us into supposing that their ideas are the same, or even their modes of thought.

It is bother enough to keep in mind the century in which a word is uttered in order to interpret its meaning. But the ambiguities are even worse, for the Old earlier meanings do not disappear; they tend to persist alongside the new. This is perhaps the most insidious part of what C. P. Snow has dubbed "the Problem of the Two Cultures." To know what a speaker means by "energy" it is not enough for me to ~~consult the calendar--to~~ know what century he is speaking in. ~~If I am not to misunderstand him,~~ I must also know whether his talk belongs to the common culture or to the scientific culture. If the former, I should not attribute to his words the quantitative precision that belongs to the latter; and if the latter, I should not interpret his words vaguely or metaphorically.

This confusion of interpretations reaches its peak in social science discourse, for historically--and ethnologically, one might say--the social sciences sit squarely athwart the boundary between the common culture (including the humanities), on the one hand, and the scientific culture, on the other. Knowing that a man is a card-carrying behavioral or social scientist does not tell you which culture he belongs to--hence does not tell you how to interpret his discourse. He may be operating in the allusive, metaphorical world of the common culture; or he may have invested his words--or tried to--with new scientific precision. I take a book titled "Value and Capital" off the shelf, and wonder what the author can mean by "value." When, on opening the book, I find it filled with mathematical equations, I conclude that "value" is not intended in one of its traditional, humanistic meanings, but in the precise scientific sense it has acquired in modern mathematical economics.

X

Old Words in New Meanings

All of this is preliminary to raising a difficulty we must hurdle if we are to communicate ~~tonight~~ today. In my talk, I intend to use familiar words like "information," "thinking," "organization." Unless I warned you, you might have every reason to suppose that these words were being used with the meanings that the common culture has attached to them over the past centuries. But if you make that--quite reasonable--assumption, you will misunderstand what I am going to say.

During the past quarter century, words like "information," "thinking," and "organization" have begun to acquire new scientific meanings, increasingly precise and increasingly quantitative. Words associated with the generation and conversion of information are today undergoing a change of meaning as drastic as the metamorphosis experienced by words associated with the generation and conversion of energy in the 18th and 19th centuries.

If you stay within the common culture, you cannot carry on a 20th century conversation about energy with a physicist or engineer. If you stay within the common culture, you will find it increasingly difficult to carry on a 20th century conversation about information with a social scientist who belongs to the scientific rather than the humanistic subculture of his discipline. The difficulty ~~in communicating~~ will not stem from jargon; it will stem from a complete disparity of meanings hidden behind a superficially common language.

What do I mean when I say: "Machines think"? The word "machine" seems obvious enough--the reference is to a modern electronic, digital computer. But "machine" has all sorts of humanistic overlays that I did not intend in my sentence. A machine, in the common culture, moves repetitively, monotonously. It requires direction from outside. It is inflexible. With the slightest component failure or mismanagement, it degenerates into "senseless" or "random" behavior.

Under certain circumstances (with certain kinds of programs stored in its memory) a computer exhibits none of these "mechanical" properties of a machine in its behavior. While retaining the word "machine" in the scientific culture as a label for a computer, I have revised drastically the associations stored with the word in my memory. Unless you are aware of that revision, you will suppose that my sentence means: "Devices that behave repetitively, inflexibly, that require outside guidance, that often become random, think."--a patent absurdity.

The word "think" itself is even more troublesome. In the common culture it denotes an unanalysed, partly "intuitive," partly subconscious and unconscious, sometimes "creative" set of processes, occurring in the human mind, which sometimes allows humans to solve problems or make decisions or design something. What does that process have in common with the processes computers follow when they execute their programs?

The common culture finds almost nothing in common between them--among other reasons, because it has never described the first-mentioned process, human thinking, but only named and labelled it. Contemporary psychological research has been discovering what is involved in the human information processing called thinking, and has been writing programs for digital computers that duplicate that processing in considerable detail. When a psychologist who has been steeped in this new scientific culture says "Machines think," he has in mind, of course, the behavior of computers governed by programs of this kind. He means something quite definite and quite precise, and something that has no satisfactory translation into the language of the common culture. If you wish to converse with him--which you well may not!--you will have to follow him into the scientific culture.

In the course of the Industrial Revolution, new groups of professionals sprang up to handle the new meanings of "work" and "energy" and the technologies associated with those new meanings: engineers, scientists, mechanics, technicians of all sorts. But the core of life in human organizations, and especially in the executive suites of organizations, was remarkably untouched by the energy revolution. The generation and conversion of energy were tasks for the blue-collar parts of organizations. Managers--except the foremen and managers whose stations were on the factory floor--dealt with the technology only remotely, abstractly, through symbols. Management activity was still mostly a part of the common culture, carried on in the language of the common culture. Only in the pecuniary calculus of the balance sheet and the ledger, at the opposite pole from the domain of "real work", did management show important signs of developing new concepts and new meanings for words.

We now witness new meanings and scientific concepts growing up about the word "information." As the science of information processing continues its development, it will not be as easy to sequester it from the main stream of managerial activity--or, human social activity--as it was to isolate the physical sciences and their associated technologies. Information processing is at the heart of executive activity, indeed, at the heart of all social interaction. More and more, all of us are finding occasion to use terms like "information," "thinking," "memory," "decision making" with 20th century scientific precision. The language of the scientific culture occupies more and more of the domain previously reserved to the common culture.

Make no mistake about the significance of this change in language. It is, as I have insisted, a change in thought, in concepts. And in this instance, it is a change of the most fundamental kind in Man's thinking about his own processes--about himself.

the of Attention Scarcity

My title speaks of "an information-rich world." How long has the world been rich in information, and what are the consequences of its prosperity--if that is what it is?

Last Easter, my neighbors bought their daughter a pair of rabbits. Whether by intent or accident, one was male, one female, and we now live in a rabbit-rich world. Persons less fond than I am of rabbits might even describe it as a rabbit-overpopulated world. Whether a world is rich or poor in rabbits is a relative matter. Since food is essential for biological populations, we might judge the world as rabbit-rich or rabbit-poor by relating the number of rabbits to the amount of lettuce and grass (and garden flowers) available for them to eat. A rabbit-rich world is a lettuce-poor world, and vice versa.

~~Once we look at the matter in this way, we recognize that~~ The obverse of a population problem is a scarcity problem--hence a resource allocation problem. There is only so much lettuce to go around, and it will have to be allocated somehow among the rabbits.

Now when we speak of an information-rich world, we may expect, analogically, that the wealth of information means a dearth of something else--a scarcity of whatever it is that information consumes. What information consumes is rather obvious: it consumes the attention of its recipients. Hence a wealth of information creates a poverty of attention, and a need to allocate that attention efficiently among the overabundance of information sources that might consume it.

To formulate an allocation problem properly, ways must be found to measure the quantities of the scarce resource; and these quantities must not be expandable at will. By now, all of us have heard of "bits of information"--a unit of information introduced by Shannon in connection with problems in the design of communication systems. We might consider measuring an information processing system's capacity for attention in bits, somehow.

Unfortunately, the bit is not the right unit. Its defects are rather too technical for detailed discussion here; but, roughly, the trouble is that the capacity of any device (or person) for receiving information, measured in bits, depends entirely upon how the information is encoded. Bit capacity is not an invariant, hence is an unsuitable measure of the scarcity of attention.

However, there is a relatively straightforward solution to the measurement problem. We can measure how much scarce resource is consumed by a message by noting how much time the recipient spends on it. Human beings, and contemporary computers also, are essentially serial, one-thing-at-a-time ~~systems.~~ devices. If they attend to one thing, they cannot, simultaneously, attend to another. That is just another way of saying that attention is scarce.

~~Last month,~~ In his paper, John Kemeny ~~told you about~~ described modern time-sharing ~~computing~~ systems, which seem able to attend to one hundred things at once[5]. They really don't, of course. Rather, they share their time and attention among these hundred things. Hence, the attention-capacity measure I am proposing will work for time-sharing systems as well as for more conventional computing systems or human beings. An organization employing many people can also be viewed as a time-sharing system, and its attention-allocation problem treated accordingly.

Scarcity of attention in an information-rich world will be measured by the time, in minutes or in hours, say, of a human executive. If we wish to be more precise, we can define a standard executive--one with an I.Q. of 120, say, a college Bachelor's degree, and so on--and ask Lew Branscomb to ~~stone~~ examine him in the Bureau of Standards here in Washington. That degree of precision won't be necessary ~~tonight~~ today. Further, we can work out a rough conversion between the attention units for human executives and for various kinds of computers.

In an information-rich world, most of the cost of information is the cost incurred by the recipient. It is not enough to know how much it costs to produce and transmit it; we must also know how much it costs, in terms of scarce attention, to receive it. I have tried bringing this argument home to my friends by suggesting that they recalculate, on this basis, how much the New York Times (or the Washington Post, if you prefer) costs them--including the reading costs. After making the calculation, they usually have exhibited alarm, but haven't cancelled their subscriptions. Perhaps the benefits still outweigh the costs.

Having explained what I mean by an information-rich world, I am now ready to tackle the main question: how we can design organizations, business firms and government agencies, to operate effectively in such a world; how we can arrange to conserve and allocate effectively their scarce attention.

I shall proceed with the help of three examples, each illustrating a major aspect of the organization problem. I have made no attempt to cover all the significant problem areas--the hour is too late for that--but have simply selected three that are sufficiently different to give a picture of how one uses this approach to organization design over a range of circumstances.

And one final caveat: any fancied resemblances of the hypothetical organizations I shall talk about to real organizations, living or dead, in the city of Washington, are illusory, fortuitous, and the product of the purest happenstance.

Information Overload

Many proposals for eliminating "information overload" (another phrase to describe life in an information-rich world) contemplate installing a computing system. There is good precedent for this. The Hollerith punched card is a creative product of the Census Bureau's first bout with information overload. A series of crises in the central exchanges of the phone company led to the invention of automatic switching systems. (In New York, at least, those particular crises don't seem to be over yet.)

Today, for example, it is argued by some that a modern postal service is doomed to collapse from information overload unless means are found to automate its sorting operations. This cannot be so. There is no reason why mail sorting costs should increase more than proportionally with the volume of mail. Hence, unit costs should, at worst, remain almost constant with volume. The cause of this particular Post Office problem (there are others, I am sure) is quite simple. If we tell someone he can have certain information processing services free, or almost free, he may demand almost an infinite amount of them. Then we find that we are not really prepared to provide the subsidy we committed ourselves to implicitly when we promised the "free" services, and we renege by performing the services badly, with insufficient resources. The crisis in the Post Office does not call for computers--it calls for the thoroughgoing application of the price system and market mechanisms.

This is not to argue that there may not be Post Office operations, including sorting, that can be performed, now or soon, more economically automatically than by hand. That kind of technical question can be settled, within reasonable limits of error and debate, by cost-benefit analysis. But there is no magic in "automation" that allows it to resolve dilemmas posed by an unwillingness or inability of an organization to allocate and price scarce information-processing resources--whether the resources are sorting clerks or electronic devices. Free or underpriced resources are always in desperately short supply. What is sometimes alleged to be technological lag in the Post Office is really failure of nerve. I won't presume to judge whose nerve has failed.

Before we rush out to buy a computer to deal with the information overload, therefore, we had better analyse the situation a little more deeply. A computer is an information processing system of quite general capability. It can receive information, store it, operate on it in a variety of ways, and transmit it to other systems. Whether a computer will contribute to the solution of an information overload problem, or compound the problem, depends on the distribution of its own attention among these four classes of activities--listening, storing, thinking, and speaking. The general design principle can be put thus:

> An information processing subsystem (a computer, a new organization unit) will reduce the net demand on attention of the rest of the organization only if it absorbs more information, previously received by others, than it produces--if it listens and thinks more than it speaks.

I could reach a similar conclusion about the perennial shortage of computer capacity on university campuses, but that would lose me some of the few friends I still have.

To be an attention conserver for an organization, an information processing system (an IPS, as I shall call it from now on) must be an information condenser. It is conventional to begin the design of an IPS by considering the information it will supply. In an information-rich world, that is doing things backwards. The crucial question is how much information it will allow us to withhold from the attention of other parts of the system.

Basically, an IPS can perform an attention-conserving function for other systems in two ways: (1) it can receive and store information that would otherwise have to be received by those other systems, and (2) it can transform ("filter") information into an output that demands fewer hours of attention than the input information.

To illustrate these two modes of attention conservation, let me talk about some of the information needs of a nation's Foreign Office. (Since the United States has a State Department and not a Foreign Office, I am obviously talking about some other country--hence, we can focus on principles and not be distracted by irrelevant particulars.) The bulk of information that flows into the system from its environment is irrelevant to action at the time it flows in. Much of it will never be relevant, but we can't know in advance with certainty what part will and what won't. Clearly, one way to conserve Foreign Office attention is to interpose ("interface" is the verb that computer specialists would use) an IPS, human or automated or both, between environment and organization, to store the information on receipt after indexing it appropriately.

The second way in which the new IPS can conserve Foreign Office attention is to analyse the information received, draw inferences from it, summarize it, and then store the products of its analyses (properly indexed again). When the rest of the system needs information, these products can be provided, instead of the raw information.

Now, this proposal has a familiar ring about it. I have simply described, in somewhat unconventional language, the quite conventional functions of a quite conventional intelligence unit. Moreover, I have solved the information overload problem simply by adding additional information processors--I eliminated scarcity by increasing the supply of scarce resources. Any fool--with money--can do that.

But the very banality of my solution carries an important lesson: The functional requirements of an IPS--how it must be designed if it is to achieve an attention-conserving goal--are largely independent of specific hardware, automated or human. Hardware will come back into the picture when we consider the costs of performing these functions and the equipment that will perform them most cheaply--but that is a later stage of the analysis.

In several respects, the proposed solution may be far less conventional than it sounds. First of all, I described the IPS as "analysing the information received, drawing inferences from it, summarizing it." In the language of the common culture, verbs like "analyse" and "summarize" are vague terms to denote vaguely defined processes carried out in vaguely specified ways by miscellaneously selected and trained human beings.

If the new IPS is to be automated, or even partly so, then we must provide much more precise descriptions, in the language of the scientific culture, of the processes denoted by these vague terms. And even if we do not intend to automate the process, the new information processing technology will permit us to formulate the programs of the human analysts and summarizers with precision, so that we can predict reliably the relation between the system's inputs and its outputs. Hence, when we look more closely at the structure of the new IPS and how it will operate, we see that it really will not resemble a traditional intelligence unit very closely at all.

Perhaps I have also exaggerated the conventionality of the solution in a second direction. The intelligence IPS I have proposed is to be designed not to supply the Foreign Office with information, but to buffer it from the over-rich environment of information in which it swims. Information does not have to be attended to--and certainly not attended to now--just because it exists in the environment. Designing an intelligence system means deciding when information is to be gathered (much of it will be preserved indefinitely in the environment if we don't want to harvest it now), where and in what form it is to be stored, how it is to be reworked and condensed, how it is to be indexed to give access to it, and when, and on whose initiative it is to be communicated to others.

Designing an intelligence system on the principle that attention is scarce and must be preserved is quite different from designing it on the principle of "the more information the better." The Foreign Office I have been talking about thought it had a particular communications crisis a few years ago. During times when events in the world were lively, the teletypes carrying incoming dispatches frequently fell behind. The solution: replace the teletypes with line printers of much greater capacity. No one, apparently, asked about the IPS's (including, presumably, the Foreign Minister) that received and processed the teletype messages, and whether these IPS's would be ready, willing, and able to process the much larger volume of messages coming from the line printers.

(My thinking on this problem has benefitted greatly from acquaintance with the analyses that have been made over the past several years of information processing requirements in the U. S. State Department. These planning activities have been laudably free from premature obsession with automated hardware.)

Everything I have said about intelligence systems in particular can be said about management information systems--a currently popular term--in general. The proper aim of a management information system is not "to bring the manager all the information he needs," but to reorganize the manager's environment of information so as to reduce the amount of time he must devote to receiving it. Stating the problem in these two different ways leads to very different system designs.

The Need to Know

And that brings me to the question of <u>the need to know</u>--how do we go about deciding where information should be stored in an information-rich world, and who should learn about it?

Those of us who were raised during the Great Depression don't always find it easy to adapt to an Affluent Society. When we ate new potatoes, we always ate the peels. ("The best part of the potato," my mother always insisted.) Non-returnable containers, even apart from their littering propensities, seem to us symbols of intolerable waste.

Our attitudes toward information also reflect the culture of poverty. We were brought up--perhaps still are--on Abe Lincoln walking miles to borrow (and return!) a book, and reading it by firelight. Most of us are constitutionally unable to throw a bound book into the wastebasket, and have trouble enough disposing even of magazines and newspapers. Some of us are so obsessed with information that we feel compelled to read everything that falls into our hands--although the bourgeoning of the mails, mentioned earlier, is helping to cure us of that obsession.

If these attitudes were highly functional in a world of clay tablets, scribes, and the human memory; if they were at least tolerable in the world of the printing press and cable; they are completely maladapted to the world of broadcast systems and the Xerox machine.

The change in the information processing technology demands a fundamental change in the meaning we attach to the familiar verb "to know." In the common culture, "to know" meant to have stored in one's memory in such a way as to facilitate recall when appropriate. By metaphoric extension, knowing might include having access, with the skill necessary for using it, to a file or book containing information.

In the scientific culture, the whole emphasis in "knowing" shifts from the storage of information--its actual physical possession--to access--the processes of using information. It is possible to have information stored without having access to it (the name on the tip of the tongue, the lost letter in the file, the unindexed book, the uncatalogued library). It is possible to have access to information without having it stored (a computer program for calculating values of the sine function, a thermometer for taking a patient's temperature).

It is obvious that if a library holds two copies of the same book, one of them can be destroyed (let me avoid the word "burned"), or exchanged, without losing information from the system. In the language of Shannon's information theory, we say that the copies make the library <u>redundant</u>. But copies are only one of three important forms of redundancy in information. Even if the library has only one copy of each book, there will be a high degree of overlap in the information they contain. If half the titles in the Library of Congress, selected at random, were destroyed, little of the World's knowledge would have disappeared, even from the Library.

But the most important, and subtle, form of redundancy derives from the fact that the world is highly lawful. To say that certain facts are random is to say that no part of them can be predicted from any other parts--each fact is independent of the others. To say that the facts are lawful is to say that certain of them can be predicted from certain others. But in that case, after storing only a fraction of them--the fraction needed to predict the others--we still retain full knowledge of them.

That is exactly what science is all about--the process of replacing unordered masses of brute fact with tidy statements of orderly relations from which those facts can be inferred. The progress of science, far from cluttering up the world's libraries with new information, enormously increases the redundancy of those libraries by discovering the orderliness of the information already stored. With each important advance in scientific theory, we can reduce the volume of explicitly stored knowledge--destroy a few books--without losing any information whatsoever. That we make so little use of this opportunity does not deny that the opportunity exists.

Let me recite an anecdote that illustrates the point very well. We are all aware that there is a DDT problem. DDT is one of those mixed blessings that technology has given us--very lethal to noxious insects, but uncomfortably persistent and cumulatively harmful to eagles, game fish, and possibly ourselves. The practical problem: how can we enjoy the agricultural and medical benefits that the toxicity of DDT has given us without suffering the consequences of its persistence?

A distinguished chemist of my acquaintance, a specialist neither in insecticides nor biochemistry, asked himself that question. From the name of DDT, he was able to write down the approximate chemical structure of the compound, encoded in the name. In the structural formula, he could recognize, on general theoretical principles, the component radicals that accounted for its toxicity. The formula also told him, on theoretical grounds, why the substance was persistent--why the molecule did not decompose readily or rapidly. He asked, again on theoretical grounds, what a compound would look like that had the toxicity of DDT, but would decompose readily. He was able to write down the formula, and could see no theoretical reason why the compound could not easily be produced. (All of this cost him ten minutes.)

A phone call to an expert in the field confirmed all of his conjectures. The new compound he had "invented" was a well-known insecticide, which had been available commercially before DDT. It was not as lethal as DDT over quite as broad a band of organisms, but nearly so, and it decomposed fairly readily. Now I don't know if the new-old chemical "solves" the DDT problem. The durability of DDT was one of the very qualities its inventors were trying to produce--so that frequent respraying would not be needed, and the costs of treatment correspondingly reduced. There may be other economic issues involved, and even chemical and biological ones.

My point is quite different. What the story illustrates is that good problem-solving capacities combined with powerful, but compact, theories may take the place of whole shelves of reference books. (And, I might add, that a telephone call may take the place of keeping the reference books, even if they are required, on your own shelves.) It may often be more efficient to leave information in the library of Nature, to be extracted by experiment or observation when needed, than to mine it and stockpile it in Man's libraries--where the retrieval costs may be as high as the costs of recreating it from new experiments or deriving it from theory.

These considerations temper my enthusiasm for using the new automated data processing technology that is becoming available largely as a means of storing and retrieving bodies of data larger than any we have ever stored and retrieved before. I do not mean to express a blanket disapproval of all proposals to improve the world's stores of information. I will even propose, a little later, that we experiment more vigorously and imaginatively than we have been willing to do thus far with the development of appropriate data banks of information about social phenomena.

What I am arguing for is the design of IPS's that will have data analysis capabilities able to keep up with our propensities to store vast bodies of data. Our capacity to analyse data will progress at an adequate pace only if we are willing to invest substantial funds--comparable to those we are now willing to invest, say, in high-energy physics--in basic research and development of computer "software," particularly in those areas that go under the headings of information retrieval, natural language processing, and artificial intelligence.

Today, computers are most often regarded as moronic robots--and that is what they are and will be as long as the programming art remains in its present primitive state. Moronic robots can sop up and store vast quantities of information, and they can spew out vast quantities. If they are to exercise due respect for the scarce attention of the recipients of that information--and that is what they are going to have to do for us--then they need to behave at a higher level of intelligence. It will take a large and vigorous research and development effort to bring that about.

In a knowledge-rich world, progress does not lie in the direction of reading information faster, writing it faster, and storing more of it. Progress lies in the direction of extracting and exploiting the patterns of the world--its redundancy--so that far less information needs to be read, written, or stored. And that progress will depend on the IPS's, human or automated, that we are able to devise for analysing and recoding information--

in brief, on our ability to devise better and more powerful thinking programs for man and machine.

Technology Assessment

Attention is *generally* scarce in organizations. It is *particularly* scarce at the tops of organizations. And it is *desperately* scarce at the top of the organization we call the United States Government. There is only one President, and although he can be assisted in certain respects by his Budget Bureau, his Office of Science and Technology, and the other elements of the Executive Office, a frightening array of matters converges on that single, serial information processing system, the President of the United States.

There is only one Congress of the United States. To a considerable extent, it can turn itself into a parallel organization by operating through its committees. But every important matter must occupy the attention of many members of that body, and highly important matters may claim substantial time and attention from all.

There is only one body of citizens in the United States. For many purposes, they can go about their several affairs, but large public problems--the Viet Nam War, civil rights, student unrest, the cities, environmental quality, to mention five near the top of the current agenda--require periodically a synchrony of public attention. And even the short list I gave above (not even mentioning full employment, the Bomb, and so on and on) is more than enough to crowd the agenda to the point of unworkability or inaction.

Let me avoid the task of describing how the entire agenda should be set--which staggers me--and focus on just one small segment of it. Our

~~chairman~~ moderator this evening, Congressman Daddario, has devoted a great deal of thought in recent years to improving the procedures in our society and Government for dealing with the new technology that we produce so prodigiously. At the request of his House Subcommittee on Science, Research and Development, a panel, on which I served, of the National Academy of Sciences recently prepared a report on technology assessment [4].

Parenthetically, I was more than a little appalled when the *New York Times* carried the story about the technology assessment report under the headline, "Technology Peril Stirs Scientists," and followed that headline with a citation of all of the real or imagined adverse effects of technology. If you will trouble to go behind these headlines to the actual report, you will find that the panel did not fall into the vulgar error, so fashionable at the moment, of blaming all our ills on technology. If the problem were only that simple! If we knew that new technology was bad, we could simply ban it.

Technology assessment is not just a matter of determining what good and bad effects new technological developments are likely to have. Even less is it a matter of making sure, before new technology is licensed, that it will have no undesirable effects. The dream of thinking everything out before we act; of making certain that we have all the facts, know all the consequences, is a sick Hamlet's dream. It is the dream of someone who has no appreciation of the seamless web of causation in the world, the limits of human thinking, or--our topic tonight--the scarcity of human attention.

The world outside is itself the greatest storehouse of knowledge. Human reason, drawing upon the pattern and redundancy of nature, can predict some of the consequences of human action. But the world will always remain the largest laboratory, the largest information store, from which we will learn the outcomes, good and bad, of what we have done. Of course it is costly to learn from experience. But it is also costly to carry out research and analysis to anticipate experience. Knowledge from the laboratory is not always cheaper--and frequently is much less reliable--than knowledge from life.

Technology assessment is an intelligence function. If it operated perfectly, which it is certain not to, it would do two things for us: First, it would warn us in advance--before we had taken action--of the really dangerous, and especially the irreversibly dangerous consequences that might flow from proposed innovations. Second, it would give us early warning--before major irreversible damage had been done--of unanticipated consequences of innovations as they became visible in nature's laboratory after they had been introduced. In performing both of these functions, it would be mindful of the precious scarcity of attention. It would only put items on the agenda on a presumption that some attention and action was needed--in some cases, only the action of gathering information to evaluate the need for further attention.

A phrase like "technology assessment" conjures up a picture of scientific competence and objectivity, of deliberateness and thoughtfulness, of concern for the long run, of a "systems" view that considers all aspects, all consequences. Those are all desirable qualities of a decision making system, but they are not qualities we can impose without considering the organizational and political environment of that system.

As our scientific and engineering knowledge grows, so does our power to adopt measures and take actions that have consequences ramifying over vast reaches of space and time. The growth of knowledge not only makes our actions more consequential, but it also allows us to recognize their consequences, where we would have been ignorant or ignored them before. We are able to make bigger and bigger waves, and at the same time we have more and more sensitive instruments to detect the rocking of the boat. Today, we sterilize and quarantine everything that travels between Earth and Moon. Less than five hundred years ago, we diffused ~~T.B.~~ tuberculosis, smallpox, and syphilis throughout the Americas in happy ignorance.

The injunction to "take account of all the effects" only conjures up the picture of an integral that stretches out through all space and time and doesn't converge. We must assume, as mankind has always assumed, that if we do a reasonable job of allocating our limited attention and our limited powers of thought, we will solve the crucial problems that face us at least as fast as new crucial problems arise. If that assumption is wrong, there is no help for us. If it is right, then technology assessment becomes part and parcel of the task of setting society's and our Government's agenda.

To bring the notion of technology assessment out of the realm of abstraction, let me go back to a specific example mentioned earlier--DDT. As far as I know--though I haven't researched the history of the matter--DDT was introduced on a large scale without thorough, or at least adequate, study of the potential dangers of its cumulation in the atmosphere and in organisms--especially predators. It was hailed, for its agricultural and medical benefits, as one of technology's "miracles." Now, some decades later, we learn that it is a somewhat flawed miracle.

The possible adverse effects of DDT have been known to specialists for some time. Perhaps they were even known, but ignored, at the time DDT was originally introduced. If that were true, it would simply underscore my fundamental theme of the scarcity of attention. Suppose it is not true--that the dangers of DDT have been discovered only in the laboratory of nature.

Now I must say something that may enrage all the eagle lovers in this audience. I am not sure that we--or even the eagles--have suffered unconscionable or irreversible loss by letting experience, actual use, tell us about DDT, rather than trying to anticipate and predict that experience in advance. The technology assessment has been made--is being made--by the environment; we are getting signals from the environment calling attention to some of its findings; and these signals are strong enough to deserve and get our attention. As you know from reading the newspapers, the DDT issue has been claiming attention intermittently for some months. The loudest environmental signal attracting that attention has been the detection of DDT in Great Lakes game fish. The issue finds itself high on the agenda of newspapers, courts, and committees, and attention may persist long enough to bring action.

I know this sounds complacent, and I really do not feel complacent. But it serves no useful social purpose to treat with anguish and hand-wringing every public problem which, by hindsight, might have been avoided if we had been able to afford the luxury of more foresight. Now that we <u>know</u> that DDT creates problems, it is more profitable to address ourselves to those problems than to hold inquests to discover who should have seen the problems earlier.

Carrying the DDT problem out of the past, into present and future, what is involved, concretely, in assessing this particular piece of technology? Our information about the effects of DDT--particularly about its effects on other animals and plants, and the effects of long-continued diffuse contamination--is in many respects quite unsatisfactory. That is true, by the way, of our information about almost any issue of public policy you may want to investigate. But to say that our information is unsatisfactory does not mean that we could improve the situation by massive collection of data. On the contrary, what we mainly need is carefully-aimed, high quality biological investigation of the cause and effect mechanisms involved in both the diffusion and metabolism of DDT. After we understand better the chemistry and biology of the problem, perhaps we could make sense of masses of data we might gather--but probably we would then not need much data.

But first-rate biologists and chemists capable of doing the research are in as short supply as most other high-quality information processing systems. Their attention is an exceedingly scarce commodity, and we are unlikely to capture much of it soon. The practical question, as always, is how to deal with the situation on the basis of the scrappy, inadequate data we have now.

We begin to ask questions like these: Assuming the worst possible case for the harmful effects of DDT, what might be the magnitude of these effects, humanly, economically, ecologically, and to what extent do they appear irreversible? What would it cost us to do without DDT entirely, humanly, economically, ecologically? What is the next-best alternative if we couldn't use it? And so on.

Now my first comment on these questions is that they are common-sense questions. You don't have to know anything about the technology to ask them--though you might learn something about the technology from the answers. My second comment is that the most effective IPS for getting answers to the questions consists of a telephone, a Xerox machine (to copy some of the documents your telephone correspondents refer you to), and a couple of very bright professionals (though not necessarily specialists) who do know something about the technology. With this retrieval system, you can extract just about anything in the world that is now known about the problem, and do it in a few man-weeks of work. (The time required will go up considerably if institutional customs trap you into holding "hearings" and "briefings," or organizing a "research project," but it would take me too far afield to pursue that issue.)

My third comment is that there are any number of locations, inside and outside the Federal government where the questions may be asked. They may be asked by a group working in the Office of Science and Technology, in the National Academy of Sciences or of Engineering, in a non-profit corporation like the RAND Corporation or Resources for the Future, in a Congressional committee. (An excellent example of the latter is the recent series of reports on steam-powered automobiles.)

The location of the investigating group is significant from only one standpoint--but that may be a crucial one. Where the group is located may determine the attention it can command, and, related to that, the legitimacy that will be accorded its findings. Attention and legitimacy are interdependent, but by no means identical, matters.

In technology assessment, legitimacy may sometimes be achieved, and even attention secured, by using the usual social credentials of science--the right degrees, professional posts, scientific reputations. But we all know of impeccable reports that are buried (perhaps some of our own), and we know of other reports that gain a place high on the agenda <u>without</u> the proper credentials. The Ralph Naders of the world show us that writing and speaking forcefully, understanding the mass media--and being usually right about the facts--can compensate for all kinds of missing union cards and lack of access to organizational channels. Rachel Carson showed that even literary excellence is sometimes enough to turn the trick.

I am strongly in accord with Congressman Daddario's view that we can and should strengthen the processes of technology assessment in our country, and make them more effective. While that strengthening will not permit us to dispense with the world itself as a major laboratory, it probably will permit us, to a modest extent, to substitute foresight for hindsight. It may have been unnecessary to have waited until all Los Angeles wept before doing anything about automobile engine exhausts. Well-financed institutions for technology assessment would, even today, probably be spending ten times as much as we are actually spending--a hundred million dollars a year instead of ten--to find out whether the steam automobile offers a long-term solution to a smog problem which we are now treating with temporary expedients.

Strengthening technology assessment means improving our procedures for setting the public agenda. It does not mean pressing more information and more problems on an already burdened President, Congress, and public. To paraphrase the harrassed householder of the story, "New problems? We already have problems ~~that~~ we haven't begun to use yet." There is no special virtue, in an information-rich world, in prematurely early warnings. We can best afford to let the world store the information for us until the time has come for us to focus our attention and thought on it.

Assessing Information Processing Technology

The final issue I should like to address tonight is itself a problem in technology assessment. We are confronted with a major new science and technology of information processing. It is only a quarter century old, and we have only the faintest glimmerings of what it will look like when another quarter century has gone by. How shall we assess it, and how shall we make sure that it will develop in socially beneficial ways?

The most visible, and superficially spectacular part of the new technology is its hardware--the computers, typewriter consoles, cathode ray tubes and associated gadgets it brings with it. These devices give us powerful new ways for recording, storing, processing, and writing information to improve and replace the human IPS's with which we have had to make do through all these generations of Man's history.

In my remarks tonight, I have tried to make two main points about the hardware of the new IPS's. First, by itself, the hardware does not solve any of our organizational problems, our problems of attention scarcity. Second, the hardware boxes will only begin to make significant contributions toward solving these problems as we begin to understand information processing systems well enough to conceive sophisticated programs for them--programs that will permit them to think in the ways that man does, and at least as well.

But each step we take toward increasing our sophistication and scientific knowledge about automated IPS's increases at the same time, and in about the same measure, our sophistication and scientific knowledge about the human IPS, about Man's thought processes. What we are acquiring with the new computers and new gadgets is something of far more basic significance--a science of human thinking and human organization.

There is no reason to suppose that the armchair will be any more effective as a scientific instrument for understanding this new technology than it has been for understanding previous technologies as they emerged. If we are to understand information processing, we are going to have to study it in the laboratory of nature. We are going to have to construct, and program and operate many kinds of information processing systems, in order to see what they do and how they perform.

The systems we construct initially (and most of those we have constructed in the past twenty-five years) will perform in all sorts of unexpected ways--most of them stupid --and will, by hindsight, seem incredibly crude. They will never pass a cost-effectiveness test on their operating performance; we shall have to write them off as ~~a~~ research and development ~~expense~~ efforts.

From the behavior of some of these experimental systems, we may learn that the new technology contains dangers as well as promises. Already, for example, there has been expressed considerable concern about the threats to privacy that the new technology might create. All of these concerns will be mere armchair speculations until we have a broad base of experience against which to test them.

Very early in the computer era, I advised several business firms, not to acquire computers until they knew exactly how they would use them, and how they would pay their way. I soon realized that was bad advice. Initially, computers pay their way by educating large numbers of us about computers. They are now, and will continue to be, the principal instruments in aiding us to replace the vague, inadequate common-culture meanings that now inhabit the words in our information processing vocabulary by the sharp, rich scientific meanings those words will have to have in the future.

All of this points, I think, to a rather clear public policy for understanding and assessing the new technology. We need public support for research and development efforts on a scale much larger than the scale on which they are supported today. These efforts should be as varied in nature as possible. They should certainly include network experiments of the sort that John Kemeny envisages[5]. They should include data bank experiments. And above all, they should include experiments in robotry, large-scale memory organization, and artificial intelligence of many sorts that will permit the building of the basic foundations of a science of information processing.

If past experience, on a very modest scale, is a fair indication, a program pursued in the experimental spirit I have indicated, will not lack by-products of great practical value. Not many of you would recognize, perhaps, the term "list processing." It denotes an esoteric development of computer programming languages motivated initially--about fifteen years ago--by pure research interests in artificial intelligence. Today, list processing concepts are deeply imbedded in the design of all of the large programming and operating systems that underlie the everyday bread-and-butter uses of computers for accounting or engineering computation.

The exploration of the Moon is a great adventure, and after the Moon, there are other objects still further out in space. But Man's inner space, his mind, has been less well known than the space of the planets. It is time that we establish a National policy to explore that space vigorously, and that we establish goals, time-tables, and budgets. Will you think me whimsical or impractical if I propose, as one of those goals, a world champion chess playing computer program by 1975, and as another, an order of magnitude increase by 1980 in the speed with which a human being can learn a difficult school subject--say, a foreign language or arithmetic?

And, ~~by the way,~~ what credibility ~~did you accord the~~ moonshot ~~goals when they were first~~ announced in 1961?

if we are willing to dedicate ourselves to material goals of this kind — if you don't like my two specific goals, substitute your own — if we are willing to set deadlines for them, if we are willing to commit resources to them, as we have committed resources to the exploration of outer space, I think that we ~~will~~ soon will have an understanding both of those information processors we call computers and of that very important and fascinating information processor we call man. And that understanding will enable us to build for the future organizations far more effective than those we have been able to build and operate in the past and present.