Noted expert on organization behavior, Herbert Simon, shares some ideas for getting information down to a system.

Information 101:
It's Not What You Know, It's How You Know It

Herbert A. Simon

We are constantly reminded that we live in an information society—more than a society, an information world. We are told that we are going to drive on a superhighway in this world; or perhaps we will stay at home while the information that is on the superhighway races toward us at the speed of light. Whoever or whichever does the moving, ourselves or the information, we are going to be immersed in a great bath of up-to-the-moment knowledge about the world in which we make our management decisions.

I'd like to offer a few cautions and suggestions as to how we may deal with the abundance of information that is promised—or threatened. Superhighways, after all, frequently offer a superabundance of traffic, as well as superscarcity of parking spaces.

When we undertake the design of any complex system—be they management, scientific, or personal and family information systems—our first task is to determine what functions, or goals, the system is to serve. (I assume that the function of any information system is to provide knowledge that will be useful in making decisions in some domain.) With every information system is associated a decision system.

The second task in designing a complex system is to identify the scarce factors: the bottlenecks that are going to place effective limits on the system's power to accomplish its goals. In the design of complex things, we cannot optimize over everything—our bounded rationality does not permit it. We must focus on those particular variables that are going to affect system performance most strongly and critically. These are the scarce factors.

Think back to the Easter holidays. Some parents buy a rabbit as a pet for their children. Since the rabbit can also help keep the lawn trimmed, a superabundance of grass may be reduced by rabbits. Cautious parents acquire one rabbit; incautious parents sometimes acquire a pair.

We all know what happens next: soon, we have not a pair of rabbits but a lot of rabbits, and not a superabundance of grass, but very little of it. The factor that was orig-

inally scarce—rabbits—has become overabundant, and the factor that was originally overabundant—grass—has become rare.

Historically, information has been a scarce factor in the decision-making process. We have often had to make our decisions without the knowledge of some of the most critical conditions that could determine outcomes. Information is precious to us. We have a horror of burning books, and many of us have a hard time throwing away papers.

Information, however, is not the scarce factor. Many of us have not yet noticed that a number of years ago we entered an era in which information was no longer scarce, in which information was thrust at us in a steadily increasing flow. But at this time, we didn't always have the information that would be relevant to our decisions (such as, whether the stock market would go up or down the next day or whether next month's orders would increase or decrease).

We were provided with large quantities of information that we could sift and filter through to find the relevant nuggets but which we could surely not absorb in their totality. If one factor in a system (in this case, information) has become abundant, what has become scarce? Our time—your time and mine—to process that information.

But who is going to filter the increased volume of information that faces us? Filtering information requires understanding both the information and the needs of the user, so as to determine what is relevant and irrelevant. In other words, it requires intelligence—either human or artificial intelligence. And as the human intelligence is already fully engaged, which is why there is a problem, artificial intelligence is what is required.

The information that people use in making decisions, especially the really important decisions such as those we often call "policy," is predominately nonnumerical. When I speak of artificial intelligence, I am not referring to the well-known ability of computers to crunch vast quantities of numbers—however important that ability may be for engineering design, scheduling, or some kinds of modeling and fiscal planning. Computers have no special affinity to numbers; they process symbols, patterns of any kind, including verbal, diagrammatic, and pictorial, as well as numerical.

Much of the information that is needed for management decisions takes the form of words (in trade magazines, correspondence, reference manuals) or diagrams and pictures (product designs, architectural drawings, and so on) and is not numerical. In designing a communication system, including its databases and artificially intelligent components, these nonnumerical forms of information will undoubtedly play a much larger role than the numerical.

Much—usually most for that matter—of the information that is important for management decisions originates outside a company. Therefore, a company's information system must be designed to include this external information, and decisions have to be made about when to generate it as a corporate effort and when to buy it.

But how can a machine decide what information I need to do my job? Machines are already making such decisions every day for many kinds of jobs. Such computer programs are often called "expert systems"; and today, hundreds—perhaps thousands—of such systems are in use. Moreover, that number is growing by leaps and bounds.

MANY OF US HAVE NOT YET NOTICED THAT A NUMBER OF YEARS AGO WE ENTERED AN ERA IN WHICH INFORMATION WAS NO LONGER SCARCE.

Let me cite some examples, not all of them dealing with top management decisions but all of them relevant to how we will be able to structure information for top management decisions.

Today it is commonplace that the information needed to determine the credit-worthiness of a prospective borrower (at least for consumer credit) can be quite adequately gathered, filtered, and evaluated by a computer program. Such decisions are made every few seconds.

Another example? A commercial product on the market that makes medical diagnoses (filtering information about symptoms) as skillfully as most internists. The list

goes on as computer programs create competent abstracts of a technical paper from just the text of the paper. Other computer programs, upon reading such an abstract, can quite competently decide whether the article should be called to the attention of an engineer who has already described (in terms of key words, for instance) his or her information needs.

Notice that such programs or sets of programs actually do more than filter information. In many cases, they not only decide what information—selected from what is accessible—should be used, but they also actually use the information to make decisions. The role of artificial intelligence need not be limited to the filtering job, it can also participate in the decision-making job.

The risks of error

But if we allow machines, however intelligent they may be, to filter the information to which we have access—and even to participate in making our decisions—how can we prevent errors from creeping into such a system? We can't. But the presence of error does not distinguish the automated or man/machine system I am proposing from the mainly human system we use now. How often does each one of us attend to information that turns out to be irrelevant (and could even be easily *predicted* to be irrelevant)?

Suppose we were to develop a scheme for evaluating our own information-gathering and information-filtering methods and habits. Just how well would we score? For example, it is rumored that a great many people read a daily newspaper—the *New York Times*, the *Wall Street Journal*, or *U.S.A. Today*. How many of the items they scan actually tell them something they did not already know (or that wasn't in yesterday's newspaper)? Of the new pieces of information they did obtain, how much couldn't be reasonably predicted? On which of these pieces of information did they act? And, if they acted at all, did it have to be in 24 hours? If you apply these tests to your own reading habits, you might conclude that you should cancel your newspaper subscriptions, sell your TV and radio, and annually purchase the *World Almanac*.

But in criticizing our present habits of information intake, I have not addressed the question of error. The issue is not whether an automated filtering system would make errors; of course it would. The issue is whether it would make more serious errors of omission and commission than the present human system does. And the answer depends on the level of intelligence we are able to build into it—a level that increases each year of our experience in building such systems.

Years ago, when flying into New York City in gray, foggy weather at La Guardia Airport, I would pray that the plane was being guided by a human pilot. Today, I pray that it is guided by an automatic landing system. Fickle affections? A switch from an allegiance to machines? Not at all, it has to do with my assessment of the state of the art—my judgment of which system is the more reliable. Perhaps 20 years ago the human system was more reliable; now the automated system is. And it is by this same criterion that we will judge information-filtering systems and automated management-decision systems.

THE ROLE OF ARTIFICIAL INTELLIGENCE CAN ALSO PARTICIPATE IN THE DECISION-MAKING JOB.

Implications for organization structure

When ocean transportation (and, to a lesser extent, land transportation) cheapened after World War II, the balance of cost between importation of goods and local manufacture shifted with a corresponding increase in international trade. We should not exaggerate the magnitude of the effect, however; for international trade levels, as a fraction of gross national product, are only a little higher than they were at the turn of the century. Perhaps more striking than the increases of international trade in goods have been increases in the international trade in investment capital (although Americans should be particularly aware that this is no new thing in their history).

The multinational corporation is probably far more a product of reductions in communication costs than in transportation costs, for it more and more decentralizes its actual manufacturing operations. In so doing, it places a new load on its communication system to preserve the reality, or the fiction that the multinational firm is still a single organization. The danger is that to maintain this reality or fiction, we will press toward centralization

of decision-making at the expense of creating and exacerbating the very information overloads that I have been warning against.

In the past, momentous questions of centralization or decentralization of organizational decision-making were often settled by the fact of physical distance. If the plant was 200 miles from the head office, the president could not wander casually into the plant manager's office to tell him or her what to do. I have met, in my career, more than one plant manager who resisted (or refused) promotion to a larger plant in the headquarters city for just this reason.

Some years ago, the wireless and the long-distance telephone began to change things in this respect. (For example, it has often been noted that the functions of ambassadors have been affected in fundamental ways.) But these are devices of narrow bandwidth and, until quite recently, considerable noisiness. A resourceful, physically remote manager could retain considerable autonomy if he or she provided the only link to the home office. Now the information superhighway, in its multitudinous forms, threatens to destroy that autonomy almost entirely.

Is centralization [illegible]?

I have been discussing the cheapening of communications, and the centralization of decision making that is its probable consequence, as though they were bad things. That is obviously too simplistic a view, but it is not a view to be ignored. We should avoid concluding that, just because it is cheap to transmit messages, communication is, therefore, nearly costless. To draw this conclusion is to commit precisely the fallacy against which I am warning: treating communications, rather than the time required to deal with them, as the scarce factor in organizational decision-making and organizational life in general.

Designing an efficient and effective communication system requires determining which parts of the communication and decision-making processes (including the filtering) are most economically and effectively done by people and which by machines. Having determined that, we can see where the bottlenecks are located and what the consequences are—in terms of the loads imposed on those bottlenecks—for centralization or decentralization.

If we have information source I at location A, source II at location B and decisions K to be carried out at A and L at B, we can begin to ask such questions as: What is the relevance of the information in I to K and to L? And similarly, what is the relevance of the information in II to K and to L? What means do we have for filtering and compacting the information in I, and carrying out partial or full analysis of it, before shipping any of it to B, and ditto for the information in II?

To carry out such analysis, we would do well to start with an analysis of the decisions themselves: the information that could and should be input into them, and the kinds of analysis that need to be performed. We need to discover also what parts of the information are in our human heads and what parts stored in various other databases.

It rapidly becomes apparent that creating an effective communication system is quite the equivalent to solving a whole system of simultaneous equations. After a few possible basic plans have been designed, they have to be evaluated by seeing to whom (or what) they assign the decision-making responsibilities, the analytic responsibilities, the data-filtering responsibilities, and so on. At the same time, it must be determined where these components of the system are best located physically. And the analysis is likely to be valid to the extent that priority is given to the scarce resource: the human ability to use the information that is directed to it.

This all may seem very obvious, but it does not correspond to the way that information and communication systems have been built in the past (and up to the present). Too often they have been designed in the spirit that information technology should, obviously, be used just because it is here. This is a treacherous basis for organizational design, and, if employed, will lead to no better results than some of the late lamented management information systems of the recent past. We can do better, but only if we keep our eyes on the scarce factor—human time and attention —and on the functions that factor can best perform in comparison with systems of artificial intelligence.

This paper was originally presented at the Carnegie Bosch Institute's "International Conference High Performance Global Corporations," April 21, 1995, in Boca Raton, Florida.

Herbert A. Simon is a professor of computer science and psychology at Carnegie Mellon University. Simon may be reached at 412-268-2787 or e-mail herb_simon@v.gp.cs.cmu.edu.