A panel of Caltech experts discusses the future of

Computers and Humanity

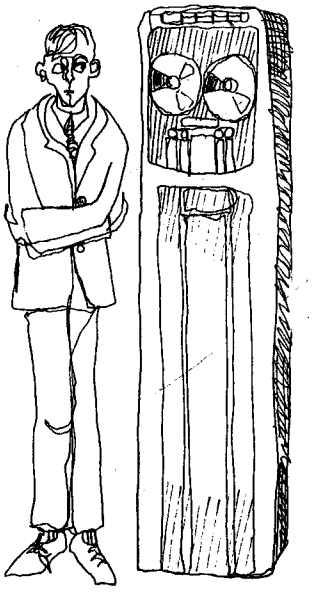

Some highlights from a panel discussion sponsored by the Caltech YMCA in Beckman Auditorium on February 10, 1965. Members of the panel were Louis T. Rader, a Caltech alumnus, who is vice president of the General Electric Company and manager of its industrial electronics division; Simon Ramo, president of The Bunker-Ramo Corp., and a member of the Caltech board of trustees; and Hallett D. Smith, chairman of the division of humanities at Caltech. G. D. McCann, director of the Booth Computing Center, served as moderator.

RADER: One of the questions most often raised is: Will computers cause unemployment?

First of all, the computer industry is one of fantastic growth. It is no more than 15 years old, but there are some 18,000 machines now in use in the United States — 10 percent of which are in the federal government. If you look at all the electrical goods that are sold to consumers — from electric light bulbs to radios, televisions, and refrigerators — then the computer industry is one-third as big as all of that. So it is a big industry already.

Some very good statistics have been accumulated in this country for the last 30 years or so by the Bureau of the Census, and I would like to cover some of them. As we look at the 20 categories into which our manufacturing is divided in this country, we can divide up these industries on the basis of those that turn out, say, a thousand dollars of goods at selling prices in terms of the man-hours required to go into them. And the four good industries are the ones that require the least labor — the chemical industry, petroleum and coal, tobacco, and instruments. These four industries will average 77 man-hours of labor per thousand dollars of output.

On the other side of the scale, the four poor industries, or the industries that require the greatest amount of labor, are textile, leather, lumber, and apparel. These industries require 270 man-hours of labor per thousand dollars of output.

You may ask: In which industries has employment grown? The first four I mentioned have the highest investment per worker also — from $30,000 per worker up to $100,000. These are the industries where we have put our money to improve the whole efficiency — with automation and computers as a part of it. And the statistics show that in the last 15 years, in the industries in which we have put the maximum amount of money for automation, employment has increased 51 percent.

The same statistics show that in the four poor industries (and I use my own definition of poor) employment has dropped 7 percent in the last 15 years. Why? Because the poor industries cannot compete with goods coming from Japan, Europe, and elsewhere. So the first thing the statistics tell us is that the industries into which we put the maximum amount of money are the industries that make

continued on page 18

the most profit, and have the greatest increase in employment.

Looking at the same figures, you say: Which industries compete best in the worldwide markets? Again, the ones that we have mechanized and automated export more than they import. The ones that are not automated import more than they export.

How do workers' salaries or wages or take-home pay compare in the industries where we put a lot of money into investment and the others? The answer is that the industries that have the greatest investment per worker pay the highest salaries.

Another way of saying this is: Is our productivity in the United States greater than Europe or Japan because we work harder? And the answer is no. Because we are smarter? No. Our productivity is greater because we give our people the tools with which to work. So, in the question of export vs. import, the question of employment, and the question of take-home pay, we can show pretty clearly that the industries in which we invest more money are the industries that do the best.

RAMO: I have been given, I believe, the favorable side of this subject. I have been asked to say a few words about the *future* impact of computers on society. When Dr. Rader talks about the situation *today,* you might challenge him, and he is in the position of having to prove that he is right. In talking about the future, you may find that everything I say seems unsound, but it is difficult for you to prove it, especially if I go far enough out in the future. And I shall – not only to make it difficult for you to ask a question that is impossible to handle, but also because I think it is essential to look out far enough, so that we can begin to appreciate some factors that are well beyond that which exists today, and which may give us the real clues as to how computers will affect society.

Since I am trying to do this in just a few minutes, I am going to do it by one example only, but I think it is a very broad one. I think it is meaningful, and typical, and substantive. But I will exaggerate a little and simplify in order to make a point quickly.

I am going to ask you to imagine that we have enough years in the past, and the use of a technology so mature and so advanced that in every aspect of our lives in which we use information, or we use our intelligence in order to achieve some useful purpose (or even just to think about useless things – just to enjoy ourselves, to produce art, to ponder and philosophize) we shall find ourselves able to do this better because we will have our brains extended by electronic devices and systems. And when I say extended, I mean we will have available the equipment of a much greater memory than is available to us today.

Picture, as this one example, what amounts to several national networks of information storage that will be continually updating the facts that apply to the professions, to education, to the running of businesses and industry, to transportation, to banking, to government. We will have the means for retrieving this information and displaying it and moving it about the country. Applied to education, this means it will be possible to introduce teaching aids that depend upon these national networks for the presentation of information. It will be possible to have statistics useful to those who are planning educational programs. In medicine, it will be possible for the individual's physician to introduce the information about his patient into the network and obtain almost instantaneously a certain kind of return information which will represent at least a fraction of what he could have gotten – in principle, at least – by consulting with thousands of other physicians who have had similar problems. This kind of thing could be used in research in disease, so as to create a relationship between cause and cure, between potentials of drugs and treatments, that could have an impact on medicine comparable with the impact of surgery.

In law, it will make possible the orderly processing of the kind of information which keeps everyone doing things according to the rules, whether it be buying and selling or forming interrelationships between corporations. In the running of companies, management will have information and, what is perhaps even more important, they will be able to create plans that are optimum, and efficient, and they will be able to create relationships between separate entities so as to assure the greatest smoothness of operation.

You might even imagine carrying this idea of

national informational intelligence one step further. The public, as a democratic public that participates in deciding issues, can be locked into the system by being asked to express their opinions – through devices that are in their homes, continually presenting issues and asking for reactions – so that the whole country can know the reaction on many issues with considerably greater competence.

All of this, of course, brings up the possibilities of informing the public, and, therefore, of having a better educated, more alert, more interested public (as well as having the possibilities of being misused, to mislead the public!).

Now this is a quick-once-over-lightly on the idea of extending the human intellect and being in some respects, then, a smarter people who can do our jobs better. Because what makes the world go round is the information that controls it. And what really determines the way in which government reacts, the public reacts, and industry is run, is information – information acquired, presented, stored, processed, and used for decision-making.

SMITH: Every new machine is a threat and a challenge. It is also, in many ways, a reward. In the seventeenth century, the telescope and the microscope were felt to be, by some people, very inhuman and antihuman. You can imagine, perhaps, the reason for this. One could say: "If God had intended us to see that far he would have made our eyes as powerful as telescopes." This is a little like the old lady who said that, "If the Lord intended us to fly in jet planes he would have had the Wright brothers invent them."

The challenge of a machine to the human sense of values is most important when the response to it is one which makes us examine our own use of our minds and abilities. And I think that the computer is a machine of this sort.

I suppose people who read that a humanist was going to appear with two scientists on this program assumed that I would oppose the computer and its possible dreadful consequences. Quite the contrary. I think that its possibilities, from the point of view of the humanities, are very great indeed. I am particularly fascinated with what has been done at the Tempo Laboratory of the General Electric Company in enabling a computer to receive and return information in a natural language – namely English. This seems to me to be an extraordinary achievement and to open up many possibilities for the further understanding, not only of the nature of language, but of the way in which we think – in words.

Now, Dr. Ramo suggested that the computer

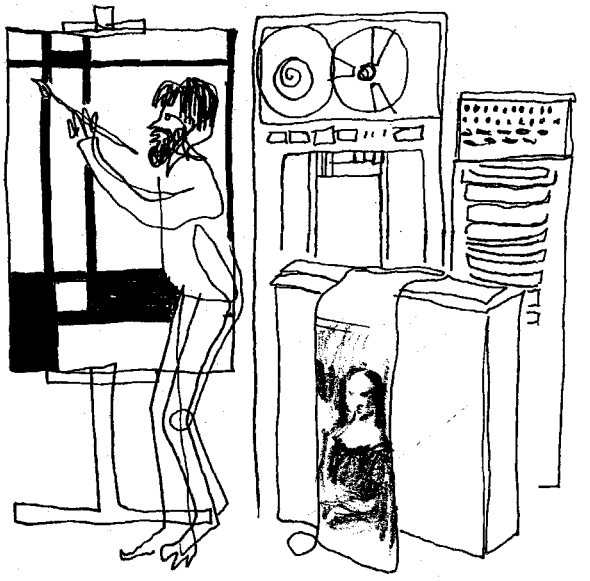

might be able to contribute something in the way of art. I am very skeptical about this. I think that, of course, you can teach a computer to reproduce any style of art. You can no doubt teach a computer to compose works in the style of Mozart. What you cannot do is to get it to create a style of its own, independent of the programming that has been put into it.

So I think there is a good deal of nonsense taught about the threat of the computer in this general area. One episode that has attracted a lot of attention in the press is the one in which a Scottish minister attempted to determine the long-debated problem of the authorship of the Epistles of the New Testament attributed to St. Paul. And, according to the papers, he put this problem on a computer and came up with the answer that five of those Epistles are by Paul, and that all the rest of them are by five other people, names undisclosed.

The sequel to this was that an American minister, who didn't like the conclusion, took the essays in which the Scottish minister had demonstrated his case, submitted them to the same test, and proved that his opponent had not written all of his own essays.

Now, what happens here, of course, is that people are using the tag of the computer for something which is really quite trivial after all. It is merely a matter of the statistics of literary style. There is a good book on that subject by an Englishman named G. Udney Yule. If you are interested in investigating, I would recommend that book to you.

But the answers that a computer gives to a question like this are no better than the material that is fed in. In this case, the statistics of literary vocabulary were inadequate to the problem.

Now, of course, people have been interested in finding out whether a computer can write poetry. And indeed it can. You feed in the words, and you program it so that they come up in any one of a number of ways, and you will get poetry. It can write beatnik poetry. When a computer writes beatnik poetry, it is called auto-beatnik, and perhaps you might be interested in a sample.

Few fingers go like narrow laughs.
An ear won't cheat jew fishes.
Who is that rose in the blind house?
And all slim, gracious, blind planes are coming.
They cry badly along a rose.
To leap is stuffy,
To crawl was tender.

I think, then, a warning is appropriate – that we should try to understand what the computer is capable of doing. And I think Dr. Ramo is quite right in saying that it will be capable of much more than we can visualize now. But we should not discount the computer because sometimes people give it trivial things to do, and because sometimes people get results from it that are nonsense.

McCANN: Before we open this for discussion, I want to add one comment, and that is that the actual principle of a digital computer, the mathematics which has so far been developed for it, really makes possible the simulation of *anything* that a human mind might conceive. Potentially, it will be possible to simulate any concept, any intellectual capacity, any creative capacity the human mind can understand.

That means a computer doesn't do just what you tell it to do; it can learn, it can teach itself – if we just knew these principles. Our big problem today, therefore, is the question not so much: What can a computer do – in breadth of intellectual activity? but: What can we learn about the principle of these activities so that we can use computers to their limit – good or bad?

QUESTION: I would like to ask what Dr. Ramo would reply to the consequences of too much leisure in society from the benefits of computers.

RAMO: I often read and hear about this concern over the relationship between the impact of computers on our society and the additional leisure time many people are going to have. I can only say, for myself, that I find it very difficult to put that anywhere near the top of my set of worries.

I find it much easier to believe that what the new electronics makes possible will so stimulate the human mind as to cause us to think up so many more things to do that it does not follow at all to me that the computer brings more leisure. It may, in effect, bring in that kind of fascinating thought and activity – impractical today and not conceived of today – that will use up more than enough time.

Take just this one thing for example: If, by proper use of electronics to aid in the educational process (and I don't mean just in schools; I mean education of the *public* on issues), we may find our public sufficiently interested in what is going on that the pondering of all the social issues will hold us busy for a good deal of the future.

RADER: Another aspect to that is that we tend to think only in terms of the standard of living in the *United States,* because we live here. There is a fantastic amount of work that has to be done, that can be done to improve the standard of living throughout the world. If we just extend our own ability, for example, in agriculture, to grow food and see to it that we find some way of developing an equivalent standard in China, Russia, India, and so on, then I think a lot of the tensions that exist in the world today will disappear. So, just because we have two cars per family, and are beginning to have two homes and so on, doesn't mean the *end* of work. We ought to have some responsibility, and I think properly *can* have some responsibility toward helping the rest of the world achieve even only a fraction of our own standard of living.

SMITH: I am not really worried about leisure. I think that there is good use of leisure and there is, of course, misuse of leisure. I suppose the sociologists would be most concerned about the rate at which this leisure became available. It is perfectly true that the industrial revolution caused by the computers is taking place at an extraordinarily fast rate, and I agree that it is a subject of some concern. I, myself, think that the notion that work is virtuous and leisure is dangerous is an inheritance from the Puritan ethic that you might just as well get over. Myself, I am in *favor* of leisure.

QUESTION: This question is directed to Dr. Rader. After the last industrial revolution we saw our number of work hours decrease from 80 hours to 40 hours per week as a result of increased productivity. So, no matter how one might advocate that more auto-

mation brings about more work, its main objective is the contrary – namely, to solve the problems of maintenance of the human society (such as clothes, food, transportation) in the minimum amount of time, to produce the leisure to free the mind for more elevated purposes.

This undoubtedly will come about as a result of automation, but it will also produce severe social dislocations in the transition period. What adaptations will our social institutions have to make so that no people go completely unemployed?

RADER: Well, there are several aspects to that question. First of all, it's an unwarranted assumption to say that increased automation creates unemployment, or that we can produce all the goods we want in, say ten hours a week instead of forty. There are no statistics at all to back up that particular statement. For example, in the last 17 years in the United States we've had an increase in employment of 13,000,000 people, and in the last 17 years we've put the maximum amount of money into our productive equipment. And the same figures are showing up in other countries of the world – Germany, France, Italy, Scandinavia, and even Japan. In fact, in some of those countries they have *over*-employment: Switzerland has to import several hundred thousand of workers to produce the goods that they want. So I don't forsee any serious economic dislocation, any massive unemployment, because all the statistics for the last 200 years say that there's more and more work as we extend our capability to do things.

RAMO: I've had occasion in the last several weeks to attempt to put down the cost, and to translate it into man-hours of employment, of updating the United States, so that it makes full use of technology in every aspect of our physical operations – transportation, banking, every form of manufacturing, even education and the professions. If you take what now appears to be both technologically and economically feasible, and if you could bring yourself to make the investment, then the payout in terms of the return on the investment would be such that the investment would be a good economic as well as engineering decision.

If you went ahead over a period of time to update the American plan and forget the rest of the world for the moment, it gets into so many trillions of dollars that it dwarfs the gross national product for any one year. The number of man-hours required are so far beyond our present total employment that it will take us many decades to get this updating accomplished – and by that time I'm sure it will be out of date.

I grant there are many problems of a political and social nature. But I just want to make it clear that, aside from these, the technological and economic aspects, put together into a calculation of what you could do that would pay off for the world, would keep everybody terribly busy creating things that we all want, could use, would enjoy having – and it would greatly raise the standard of living in the process.

QUESTION: I want to pose a kind of science-fiction question. All the things that I've read about computers that tend to make me quiver are the ideas of where computers are advanced enough so that they can actually think from original raw materials. That is, a human being is not necessary to put materials in; the computer can do everything a human being can, and is sort of self-reproductive. Now, at the present time, it seems as though the human being is the sensing element which is necessary to observe facts and put them into a thing which computers can use. Is it possible that at a future date a computer will be able to have all the sensing capabilities that we have – parallels to sight, hearing, and so forth – and will be able to take initial raw data and transpose it into usable information and thus actually think as we do, and be able to do things all by itself?

RAMO: Well, first of all, if you want to theorize about what we as human beings can conceivably and eventually cause computers to do and to be like, then you have to say that what you suggest is at least possible.

But what do we mean by the word "think"? Surely we will grant that a good bit of what we do with our minds is pretty clear and rather mundane. In data processing we have some stored facts in our minds. We put that together in certain patterns that we've learned, against information that comes in, and we arrive at certain conclusions. Now, when you understand what it is that you do with your mind and can lay it out clearly, quantitatively, or in a logical sequence, then of course you can arrange for a machine to do it. It may be, in today's technology, a complicated and impractical machine, and one that would be expensive to build and use for that purpose – since the human being is created with relatively cheap labor and costs only $20,000 or so a year to maintain. It would be foolish to set

out to duplicate it for those functions.

Then you have a more complex thing that we don't understand. We speak of it as intuition, or the thinking that our wives do, but what's really at the core of this is that we don't understand it well enough to set out to simulate it.

You have learning; you can certainly give a machine a few rules about the game of chess and let that machine play against a good chess player, and we know that in principle the computer can have a memory beyond that of a human being, even of the best chess player, so it could remember many more moves and keep track of possibilities in great detail, and eventually it could beat the best chess player.

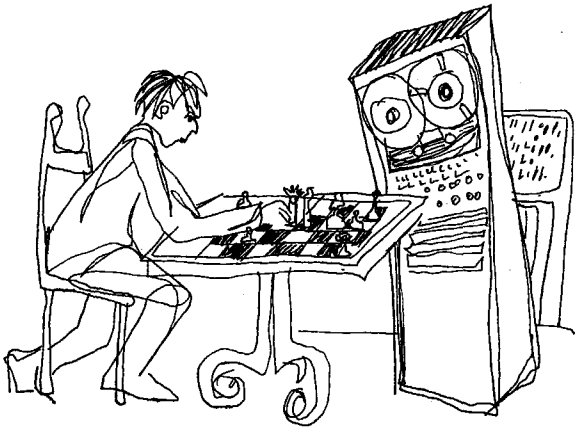

And now would the chess player argue that the machine can't think and hasn't learned how to play chess? He might so argue, but he'd be in the position of one of those who might be among the first to have his brains replaced by the computer. He's really lost his case by the actual action. So, a lot depends on your definition.

Even in the writing of music, what do we mean by the creative thought process? We don't understand it – what a composer or a true artist goes through. (I'm speaking now of the use of the computer as a tool to *help* the artist; not a replacement for the artist.) We don't really understand invention and creative thought. We don't quite know what it is that takes place.

RADER: All I'd like to say is that it's very hard for the engineers to keep up with the science writers, but we do try. There are some techniques which have derived from the transistor, from the solid state physics, which of course have changed our capability in doing things fantastically, and these techniques are constantly being extended so that we now see ways of making very large memories at low cost. We're getting closer, for example, to simulating the amount of information that can be stored in the human brain, although we're many, many degrees away from it. And we also see rather unique ways of putting together electrical circuits, so that the engineer's ability to simulate or to make a machine which can do things is progressing quite rapidly. As Si says, nothing is really impossible. So, it is possible in the far future that we can have this automatic machine that thinks – subject, of course, to the right definitions of the word think.

SMITH: With respect to the heart of the original question and its science-fiction aspects, I am not very much alarmed about the possibility of the machine sensing through various senses, and relating that to any amount of stored memory. I'm not alarmed about its ability to form concepts and to use them. I *would* be alarmed if I thought the computer could fall in love; if I thought a computer could suffer; if I thought a computer could die and understand the meaning of death, I would be very much alarmed. I don't think there's any chance of that whatever.

My favorite science-fiction fear about the computer is: You know, it's perfectly possible to get computers into this decision-making, information-gathering business so much that the computer could wage war, could press the button, collect all the information, make the decision to press the button at the right time, and so on. And presumably there would be computers on both sides. The only difficulty is that the computers would keep on fighting each other long after everybody on both sides was dead.