DEVELOPMENT OF COMMUNICATION AND COMPUTING, AND MY HOBBY

Claude E. Shannon

Thank you very much. It is a very great honor to receive this Kyoto Prize, and Mrs. Shannon and I are finding much pleasure in visiting Japan for the first time. When we began to plan our journey, Betty brought home four books relating to the culture and customs of Japan. I am not sure how reliable these were, since they were all written by Americans. One thing they had in common was telling a Japanese proverb―The nail that sticks up gets hammered down. It occurred to me that in receiving this Inamori Prize I was just such a nail, and standing up here I don't see any hammers over my head at the moment.

I don't know how history is taught here in Japan, but in the United States in my college days, most of the time was spent on the study of political leaders and wars―Caesars, Napoleons and Hitlers. I think this is totally wrong. The important people and events of history are the thinkers and innovators, the Darwins, Newtons and Beethovens whose work continues to grow in influence in a positive fashion.

While the arts and literature may be traced back for millenia, most of science goes back but a few centuries. People like Galileo, Newton and Maxwell discovered fundamental laws of the physical world which describe the motions of falling objects, of the planets, and electrical phenomena. These scientific types, of course, are still continuing with the Einsteins, Yakudas and Von Neumanns of the present century, and the effects of their discoveries tend to grow exponentially with time.

The fundamental discoveries of such scientists are wonderful achievements in themselves, but would not affect the life of the common man without the intermediate efforts of engineers and inventors―people like Edison, Bell and Marconi. Most of these inventions and applications have been developed in the last two centuries, indeed, the majority in the last century. The spinning jenny, Watt's steam engine, and the telegraph were the forerunners of this Industrial Revolution, all developed just two hundred years ago.

The telephone, electric light and radio are all about a century old now. The hundredth anniversary of the automobile was celebrated in the United States with considerable fanfare just two months ago. And of course, the airplane is considerably more recent—the Wright brothers' first flight was in 1903.

Of course it takes time, after the first success of an invention, for it to become part of daily life. Radios were not common until the 1920s, television in the 1940s. Edison's electric light, invented in 1880, required for home use the construction of power plants and a great network of power distribution lines as well as manufacturing facilities for all these components. All told, a century ago, our houses were barren of the modern everyday conveniences—no telephone, automobile, radio, television, electric light, central heating or microwave ovens. People lived much as they had centuries before, a largely agrarian society with little mobility or distant communication.

Science and technology build on themselves in an exponential way, and the technical advances of this century (or, in some cases, harmful developments) certainly exceed those of all the previous centuries combined. The technological advances of the industrial revolution can be loosely divided into three general areas—those whose ultimate aim is the processing and production of material goods, such as factories; those concerned with production and processing of energy, such as steam of solar power systems; and, finally, those related to the communication or transformation of information. The growth of this last area has in the last fifty years been even faster than that of the other two.

It has been my good fortune to be involved with many of these developments in the areas of communication and computing. Communication is basically the transmission of information from one point to another, and computing the manipulation and transformation of information.

I would like to share some of my experiences in this field, which has for me been both profession and hobby (or, more precisely, what you call "shumi").

As a student at the University of Michigan I recall reading a paper by R.V.L. Hartley which impressed me very much. It related to the transmission of information by means of various channels. During the decade after that I spent many hours on this problem, attempting to include factors such as noise in the channel and the probabilistic aspects of information. In 1948 I managed to put together a mathematical theory of communication representing my ideas up to that time. Since then my life has been largely controlled by information theory, as even now with this Kyoto Prize.

To most people the word “information” suggests meaning and reality. To the communication engineer it is the problem of getting a wave form from one point to another or, more simply, a series of letters, or, simpler still, a series of zeros and ones.

The chief concern of information theory is to discover mathematical laws governing systems designed to communicate or manipulate information. It sets up quantitative measures of information and of the capacity of various systems to transmit, store, and otherwise process information.

Some of the problems treated relate to finding the best methods of using various available communications systems, the best methods for separating the wanted information, or signal, from the extraneous information, or noise. Another problem is the setting of upper bounds on what it is possible to achieve with a given information-carrying medium (often called an information channel). While the central results are chiefly of interest to communication engineers, some of the concepts have been adopted and found useful in such fields as psychology and linguistics.

The following is the type of communication system that has been most extensively investigated:

(A) An information source that produces the raw information, or “message,” to be transmitted.
(B) A transmitter that transforms or encodes this information into a form suitable for the channel. This transformed message is called the signal.
(C) The channel on which the encoded information, or signal, is transmitted to the receiving point. During transmission the signal may be changed or distorted. The static in radio reception and the snow in television reception are familiar examples of such changes. These disturbing effects are known generally as noise.
(D) The receiver, which decodes or translates the received signal back into the original message or an approximation of it.
(E) The destination or intended recipient of the information.

Incidentally, a communication system is not unlike what is happening right here. I am the source and you are the receiver. The translator is the transmitter who is applying a complicated operation to my American message to make it suitable for Japanese ears. This transformation is difficult enough with straight factual material, but becomes vastly more difficult with jokes and double entendres. I could not resist the temptation to include a number of these to put the translator on his mettle.

Indeed, I am planning to take a tape of his translation to a second translator, and have it translated back into English. We information theorists get a lot of laughs this way.

It will be seen that this system is sufficiently general to include a wide variety of communication problems if the various elements are suitably interpreted. In radio, for example, the information source may be a person speaking into a microphone. The message is then the sound that he produces, and the transmitter is the microphone and the associated electronic equipment that changes this sound into an electronic wave, the signal. The channel is the space between the transmitting and receiving antennas, and any static or noise disturbing the signal corresponds to the noise source in the schematic diagram. The radio receiver converts the received signal into an audible output from a loudspeaker. The destination is a person listening to the message.

A basic idea in information theory is that information can be treated very much like a physical quantity, such as mass or energy. For example, an information source is like a lumber mill producing lumber at a certain point. The channel might correspond to a conveyor system for transporting the lumber to a second point. In such a situation there are two important quantities: the rate R (in cubic feet per second) at which lumber is produced at the mill and the capacity C (in cubic feet per second) of the conveyor. These two quantities determine whether or not the conveyor system will be adequate for the lumber mill. If the rate of production R is greater than the conveyor capacity C, it will certainly be impossible to transport the full output of the mill; there will not be sufficient space available. If R is less than or equal to C, it may or may not be possible, depending on whether the lumber can be packed efficiently in the conveyor. Suppose, however, that there is a sawmill at the source. This corresponds in the analogy to the encoder or transmitter. Then the lumber can be cut up into small pieces in such a way as to fill out the available capacity of the conveyor with 100 percent efficiency. Naturally, in this case a carpenter would be provided at the receiving point to fasten the pieces back together in their original form before passing them on to the consumer.

If this analogy is sound, it should be possible to set up a measure R, in suitable units, giving the rate at which information is produced by a given information source, and a second measure C that determines the capacity of a channel for transmitting information. Furthermore, the analogy would suggest that by a suitable coding or modulation system, the information can be transmitted over the channel if and only if the rate of production R is not greater than the capacity C. A key result of information theory is that it is indeed possible to set up measures R and C having this property.

If I toss a coin, it has a 50-50 chance of coming down heads or tails. This generates one binary digit or one “bit” of information. If I tossed it three times it would generate three bits or $\log_2 8$. There are eight equally likely outcomes. The general formula is $-\sum_1^{\infty} P_i \log_2 P_i$, where the P_i are the probabilities of the various events.

More complicated formulas get involved when you have correlations between events. For example, in English text, consonants and vowels tend to alternate. Channels also can become complex in various ways. For example, there can be correlations between the noise of the channel acting on successive transmitted symbols. This can lead to complexity in calculating the capacity of a channel as well as optimal encoding for it.

I would like to turn now from communication to a subject of rapidly growing economic and social importance—computers.

Computers have gone through a number of stages in terms of the components and construction. The earliest ones of Pascal and Babbage were complicated mechanical devices of great ingenuity. On the analog side, John Napier of Scotland made the marvelous invention of logarithms in 1614, and a few years later a multiplying device, Napier’s Bones, which evolved into the slide rule. The slide rule, of course, for centuries was the basic computing tool of engineers. I well remember one of my first classes in electrical engineering, when our professor told us “You’ll need to buy a slide rule for this class.” I bought a log-log-duplex, the biggest they had, and I still have it. Like most analog computers, the slide rule has become obsolete, replaced by hand held transistor computers such as this one. This does everything my log-log-duplex did and much more and out to ten decimal places instead of three…made in Japan.

In 1936 I was just graduating from the University of Michigan and wondering about a job, when it was my good luck to see a notice from the MIT Electrical Engineering Department seeking someone to operate the differential analyzer, an analog machine for solving differential equations. It was my further good fortune to obtain the job and operate the machine for the next two years. Vannevar Bush who was its chief inventor liked to call himself a “Yankee tinkerer,” but he was in fact a very sophisticated engineer—dean of engineering at MIT, and later head of the Carnegie Institution in Washington.

The differential analyzer solved differential equations up to the sixth order. It was the smartest computer of its time, but analog computers were basically doomed by the speed and precision possible in electronic and later transistor devices.

The differential analyzer, curiously, had a fairly complicated relay circuit associated with it. Relays are rather simple-minded devices. They are either off or on — the contacts are either open or closed. Contacts can be connected in series, where both must be closed to complete a connection, or in parallel where if either is closed the connection will be completed. These are close to the “and” and “or” concepts of logic and Boolean algebra, and it occurred to me that one could apply Boolean algebra to the analysis and design of relay and switching circuits. This union of Boolean algebra and switching circuits seemed to “play” as the musicians say. Using Boolean algebra one could design circuits using fewer contacts and less work. It also “played” for me, as a master’s thesis, a fellowship for a doctoral degree, and a job at Bell Telephone Laboratories.

In the 1940s several groups began working anew with digital-type computers of one type or another. Among these were Howard Aiken at Harvard University, with two relay computers which he called Mark Ⅰ and Mark Ⅱ; at Bell Telephone Laboratories George Stibitz designed computers also using relays which could carry out complicated sequences of calculations.

Early telephone relays will open or close up to a dozen contacts. In the 1940s more miniature relays were developed, and later still smaller ones.

More complex switching operations could be perfomed by special devices such as this rotary switch relay, which has eight contacts, each of which steps over 25 positions in sequence.

At about the same time, Presper Eckert and John Mauchly had constructed the ENIAC at the University of Pennsylvania, a computer using vacuum tubes rather than relays.

A vacuum tube, of course, is a much faster device than a relay, by a factor of 1000 to 1 or more, allowing computations to be done at a much higher speed. However, as time went on, the vacuum tube also became smaller as in this model. The vacuum tube had disadvantages, however, vis-a-vis the relay. It was not nearly so adept at controlling many different circuits simultaneously. It required continuous filament power and had a very limited life, measured at best in a few thousand hours. While this limited life was tolerable for a home radio with perhaps six tubes, if you build a computer with several thousand tubes, you will have a tube burn out every hour or so.

This first vacuum tube computer, the ENIAC, was used in ballistic calculations in World War Ⅱ. It was "programmed" in a sense as the differential analyzer was, by connecting various of its circuits together. It could be set up to solve a particular problem well, but could not do much in the way of making decisions.

A consultant in this work was the great mathematician from the Institute for Advanced Study in Princeton, John von Neumann, perhaps the greatest mathematician of this century, responsible for many great advances in pure mathematics, mathematical physics, and game theory. Studying the computer structure, he realized that the sequences of operations a computer does in a particular problem is itself a kind of computation — a formula akin to the formula a symbolic logician writes — and that this should not be plugged in as the ENIAC required but should be stored in memory. It should be capable of easy modification, even during computation, indeed, even itself being an object of computation.

The intellectual progress in computers in this period was so rapid that they were obsolete even before they were finished. The idea of a stored program that could be manipulated was introduced into the ENIAC's successor, the EDVAC.

These strange names of early computers were acronyms—ENIAC was Electronic Numerical Integrator and Computer; EDVAC meant Electronic Digital Vacuum-tube Computer. A few years later, as a joke, I designed a small desk computer which operated entirely in the Roman numeral system—the Is, Vs, Xs and Cs of antiquity. It was called THROBAC, an acronym for Thrifty Roman Numeral Backward-looking Computer.

In 1948, I was working at Bell Telephone Laboratories. One day I was chatting with William Shockley and noticed on his desk a small plastic object with three wires extending from it. I asked Shockley what it was and he said "It's an amplifying device like a vacuum tube, but using solid state physics." This was my first glimpse of a transistor, quite possibly the greatest invention of the 20th century. Shockley, Brattain and Bardeen won the Nobel Prize for this invention.

The transistor had tremendous advantages over the vacuum tube which it largely replaced within a few years. It was much smaller, required no filament current, and had unlimited life. Since that time, of course, transistors have replaced vacuum tubes in virtually all applications in communications and computing machines. Through the years their small size has become smaller and smaller. The tiny 1/4 inch transistor is transformed into a microchip today with hundreds of transistors inside. Just a month or so ago, Bell Laboratories announced a new even more miniaturized microchip, which can store a million bits in a chip this size. Of course, I am carrying coals to Newcastle to speak to you of microchips.

I'd like to turn now to the "shumi" side of my life with computers.

Around 1950 it occurred to me that it would be interesting to construct a machine which would solve mazes. Psychologists often use mazes as a kind of IQ test for mice. I decided that my mouse would be basically a bar magnet, moved by means of an electromagnet under the floor of the maze. The bar magnet, covered by a mouse-like shell, could be turned, and when it hit a wall of the maze could signal a computing circuit. The computer would then cause the mouse to try a different direction.

The strategy by which the machine operates can be described as follows: There are two modes of operation, which I call the "exploration strategy" and the "goal strategy." They are both quite simple. The exploration strategy is used when the mouse is first trying to find the goal. For each square in the maze, there is associated a memory, consisting of two relays. These are capable of remembering one of four possible directions: north, east, south, or west. The direction that is remembered for a square is the direction by which the mouse left the square the last time it visited that square. Those are the only data the machine remembers about the course of the mouse through the maze.

In exploration strategy, the machine takes a direction D and rotates it 90° as the first choice when it comes into a square. If it hits a barrier and comes back, it again rotates 90° and so on. When it hits the goal, a relay operates and locks in, and the machine then acts according to the goal strategy.

In the goal strategy, the machine takes as its first choice the direction by which it left the square on its last visit, and follows a direct path to the goal. Since the drive mechanisms and relay computing circuit were all under the maze floor, some of my persnickety friends complained that the mouse was not solving the maze, but the maze was solving the mouse.

The computing circuit used about 100 relays, and its trial and error procedure would solve any maze that had a solution.

This early learning machine, primitive though it was, aroused considerable interest both in the general public and among scientists interested in cybernetics and scientific aspects of learning and brain function. It was featured in the magazine *Life*.

Maze-solving mice had a revival thirty years later, when, in 1979, the Institute of Electrical and Electronic Engineers held “The Amazing Micro Mouse Maze Contest” in New York. The mice had to be self-contained—no under-the-floor electromagnets or other shenanigans. With the miniaturization of thirty years, it was possible to do this. However, the mice were the size of cats.

Incidentally, they gave me a little memento for my early work on this very important problem.

The game of chess has always been considered a high level intellectual pastime. In the early 1800s the ingenious inventor Von Kempelen demonstrated a machine that was supposed to play chess. Actually, it was something of a trick or hoax—the thinking was done by a man who was cleverly hidden inside the machine. He moved about from one place to another as the various doors were opened to convince the audience there was no man inside.

A more honest attempt to design a chess-playing machine was made in 1914 by Torres y Quevedo, who constructed a device which played an end game of king and rook against king. The machine played the side with king and rook, and would force checkmate in a few moves however its human opponent played. Since an explicit set of rules can be given for making satisfactory moves in such an end game, the problem is relatively simple, but the idea was quite advanced for that period.

In the 1940s, the computer field was advancing rapidly, and, being interested both in chess and computers, I spent some time analyzing how one might program a computer to play chess, and wrote a paper dealing with this problem. I also built, using relays, an end-game player. This picture from about 1950 shows the chessmaster Edward Lasker with me sitting at this machine. It is still very primitive by modern standards. Lasker, by the way, in addition to being a chessmaster, was a first-rate go player, and helped popularize that game in the United States with a book he wrote on the subject.

The ideas of my paper on chess-playing machines have been used and improved by many workers in this field since that time, and every year new and better chess-playing programs appear on the scene. In 1980 I attended an international tournament, in Austria, of chess-playing computers with about 20 entries. The computers being used were actually in such distant places as Canada and Italy, and were connected by telephone links. I was happy to see the tournament won by an entry called Belle from my old employer Bell Telephone Laboratories, programmed by Ken Thompson.

Incidentally, the long-time world champion chess player, Botvinnik, who is an electrical engineer by profession, has been interested in computer chess. On a visit to Russia twenty years ago, I arranged a meeting with Botvinnik, and we discussed some of the problems of machine chess. After the discussion, I challenged him to a game, to the amusement of my interpreters and colleagues. Partway through the game, I won the exchange—a rook for a knight—but inevitably Botvinnik produced some slick combinations and won with a clever mate.

Since those days, computer chips and robotics have made it possible to build cheaper and cheaper but smarter and smarter chess-players, and they have been appearing on the market by the drove. Each year the new models play stronger games, and each year I buy all the new machines.

Most of these only signal the move to be made by flashing lights, but this one indicates the move to be made by actually sliding the chesspiece to be moved by means of an electromagnet under the chessboard moved in its two dimensions by two motors—almost identical with the mechanism used for my maze-solving mouse thirty years earlier. Another interesting way of showing the moves is used by this one. It has a mechanically driven hand that moves in three dimensions, as well as opening and closing. For example, if it wishes to capture a piece, it reaches over and down, grips the piece, puts the captured piece over on the side of the board, then moves back and grips its own capturing piece and moves it to the vacated square. In the movie "The Lonely Guy," this machine appears playing a game with a human, and after winning the game shakes hands with its human opponent.

At the present time, the best chess-playing programs have reached the level of master, with ratings of over 2000. There is even an international association, with an elegant journal, devoted entirely to computer chess. I would not be surprised if, in a few more years, the world championship was held by a computer.

Recently at an Information Theory conference in England I gave a talk. The audience was falling asleep until I pulled three balls from my pocket and started juggling. Juggling skill ranges from manipulating three objects to the ten balls juggled by the great Italian, Enrico Rastelli. While my record is at best five, I began to speculate a few years ago about building machines that would either actually juggle or give that illusion.

After considerable work, I built first a machine which gives the illusion of various world record feats of juggling. Three champion jugglers are on the stage. On the right, the Romanian, Virgoaga, juggles seven clubs. They spin in the air and pass from hand to hand. On the left, Ignatov of Russia, probably today's greatest performer, juggles eleven rings. Center stage, the great Rastelli handles ten balls, five in each hand. Actually, these objects do not travel freely through the air but are supported by thin black wires from the backdrop, which are moved in their trajectories by a very complex backstage mechanism.

When I exhibited this at a juggler's convention, it aroused moderate interest, but it was clear the jugglers wanted to see real flying objects—so I spent some time building this model of W.C. Fields, the great American comedian who was also a vaudeville juggler. W.C. bounce juggles three steel balls on a drum head, and it is actual juggling, no hidden wires. He has juggled for many hours without a drop. Moreover, this is, in effect, blindfold juggling, with no feedback from the ball positions to the juggler's hands.

I believe that one could design a machine today which would establish the world record in numbers of objects—now twelve rings, held by Albert Lucas.

Building devices like chess-playing machines and juggling machines, even as a "shumi," might seem a ridiculous waste of time and money, but I think the history of science has shown that valuable consequences often proliferate from simple curiosity.

The problems one encounters in constructing such special purpose machines to emulate the human brain in such limited areas as juggling, chess playing or maze solving make one realize what a fantastic machine the human brain is. It has about 10 billion neurons. These operate in the millisecond range of speed. At birth, there is very little preprogramming apart from the basic bodily functions. The things we know and do as adults are almost all learned, from walking to language to music and mathematics.

While computers are getting smaller and smaller and faster and faster, they are still a long way from the human brain. They can make us look like idiots at long, complicated arithmetic or even logical calculations, but they cannot walk very well, and cannot recognize their own "mother." While some "learning" programs have been written, they are still very primitive. Sensory organs of robots are very inferior to the human variety, and the output mechanisms are a far cry from the many degrees of freedom of, say, the human hand, with its sensitivity to temperature and touch.

How is it that the brain outperforms these computers, even though the transistors operate in the microsecond range, a thousand times faster than neurons, which work in the millisecond area? Probably the main reason is that the neurons are working in parallel, while our computers are programmed to do but one operation at a time. When our eyes see an object, many millions of neurons are activated, sending simultaneous information to the occipital lobe of the brain. There, millions of simultaneous operations are initiated, leading to the incredibly fast recognition of faces and objects—the information is processed in parallel by the eye, not scanned and processed point by point as by a television camera.

However I have great hopes in this direction for machines that will rival or even surpass the human brain. This area, known as artificial intelligence, has been developing for some thirty or forty years. It is now taking on commercial importance. For example, within a mile of MIT, there are seven different corporations devoted to research in this area, some working on parallel processing.

It is difficult to predict the future, but it is my feeling that by 2001 AD we will have machines which can walk as well, see as well, and think as well as we do.