Lightning

BY HOWARD H. CAMPAIGNE

Secret

A discussion of techniques that may soon be put into effect to increase the speed of machine computation.

In the past, the National cryptologic effort has sometimes succeeded in the face of cryptographic competence by mounting a computing effort thought by others to be impossible. The German ENIGMA is an example, taken from a past which is sufficiently remote to be almost innocuous. Most of my readers can name other examples. So long as it is possible to have computing facilities in excess of what others may consider feasible, it behooves us to have them. Eventually we foresee that natural limitations on speed and size will be encountered, and then the inevitable advances of our opponents will corner us, so that the duel will become one of pure wits. But while we can we must maintain our superior weapons.

Project LIGHTNING was set up to preserve our advantage in speed of computation.

A word about the meaning of speed in a computer. A unit of time for a given machine will be that time in which a bit of information can be transferred from one place to another and allow the machine to recover sufficiently to begin another transfer. The number of such units in a second is the repetition rate, as "one megacycle" for one million transfers per second. Many machines have a "clock" which ticks off pulses; these pulses initiate transfers, and are far enough apart to allow the transfer to be completed and the machine to recover its balance. In such a machine the clock rate is the figure which we use for speed.

A second remark about speed of computation. It has probably occurred to the reader already that what NSA needs is not necessarily speed, but capacity. We could almost achieve our objective by having, instead of one monster n times as fast, n copies of our present machines. There are some problems for which this approach is fallacious—nine women cannot get a baby in one month—but by and large it holds for our work. There are the advantages of not putting all our eggs in the one basket, of working with machines we know, machines which can now be bought, and so on. Our answer is this: We want to have every advantage paralleling can offer, and we want the high speed, too!

Present speeds can be said to be one megacycle; this is the readily achievable production speed, not the track record. Some computers now building will come near fifty megacycles in some parts of the machine. The harsh realities of capacitance and inductance insure that our present components and circuits thereof could not do better than three hundred megacycles, and large assemblages have to move more slowly, perhaps at most sixty megacycles. In the latter aspect they act just as do automobiles. Special cars can go over 200 miles per hour, races are run at 140, but productive traffic moves very little faster than one fifth of the maximum, or 40. This drastic reduction is necessary for safety, reliability, and economy.

Since present techniques seem limited in speed we shall need new ones to win this race. The maximum of 200 miles per hour of the auto is easily bested by a jet airplane. The object of LIGHTNING is to find such new techniques. The goal is set at 1000 megacycles.

This is an ambitious goal. It is as though we, in 1925 when the speed record was under 300, had set ourselves the object of going 1000 miles per hour. We knew nothing about jet engines, nor about the sound barrier.

The program was laid out in three phases. Phase I was to be initial exploration, looking for physical phenomena which could possibly be exploited in electronic components with these speeds. Phase II was to be early development with the most promising of the results of Phase I. Phase III will be the completion of the experiments necessary to learn how to go this fast. At the end of the program it should be possible to start the construction of a computer. The kind of computer will depend on the needs at that time.

The program was undertaken by the very best talent to be found in the United States. The help of the giants in the computer field, International Business Machines, Remington-Rand Univac, and the Radio Corporation of America, was enlisted on the main problem. Massachusetts Institute of Technology was encouraged to follow up numerous long shots. Ohio State University began work on a generator of pulses at these frequencies. Kansas University invented a cathode-ray oscillograph capable of showing such pulses, for standard instruments cannot. Philco undertook to make a special component to work at these speeds.

Phase I ended in December 1958. We had found out a number of things. One was that measuring pulses at these frequencies is quite difficult. Ordinary scopes cannot follow such signals. Several sampling scopes were invented; these do not try to follow the signal but measure it at one instant. By making a number of such measurements a sampling scope can build up a fair picture of a repetitive signal.

Transient signals are beyond our art. One scope built from a traveling-wave tube is quite successful, although such a rapid scan leaves a very faint trace, hard to see. A combination traveling-wave tube and iconoscope puts out a signal which can be viewed on a television receiver. Although these various inventions better the art, they are still not fast enough. They fail to show harmonics in sufficient detail.

We learned that to be fast our components must be small and close together. In a milli-microsecond light travels about one foot, and an electronic signal will be slower. If one organ of the computer must know the result of a previous pulse in another organ in order to proceed, then the two must be less than a foot apart or the procedure will be delayed. This is a basic physical fact, and applies to whatever approach we take. Within the components the same thing is true. To make transistors work at 300 megacycles, the active region must be submicroscopic.

We found promising techniques. Let me try to say in a few one-syllable words what one approach is. Most electrical circuits have a resonant frequency, just as a bell has. If we have such a circuit, and if it is influenced by (coupled to) another circuit which is being driven at twice the frequency, our resonant circuit will begin to oscillate at its own frequency with sympathetic vibrations, just as a bell would. Its own frequency is a "subharmonic" of the driving frequency. There are two possible versions of this subharmonic, said to differ in "phase." One starts up while the other starts down; they are mirror images one of the other, or more important, one looks like the other slid sideways half a wave length, or "180 degrees." Our free-running circuit may start in either phase, we cannot tell which, but having started, it will maintain that phase. If we have a number of such circuits their phases can be compared, their phases can be controlled, and a half-wave-length of wire will reverse the phase where we want it reversed. We can thus use one phase for "zero" and the other for "one."

In fact a functioning computer has been built in Japan using phase-modulated subharmonics (but not of microwaves) so that the principle has been proved in part. There are many attractive features of this approach. One is that it seems possible to speed-up even more by using higher driving frequencies and smaller parts. Another is that many of the parts are very simple. At this time there seems to be no reason why it should not succeed. Building computing elements of such circuits is not hard; a picture of a one-digit adder is given as Fig. 1.

We found that very fast circuits may be built from "cryotrons," elements which, in order to work, depend on being at a temperature near absolute zero. At such temperatures, some metals become "super-

conductive," that is, lose all their electrical resistance. This state can be switched off with a magnetic field, so that by controlling a magnetic field we can control the resistance of a wire. The ratio of the resistances in the two cases is infinite (so far as we can measure it) so that the switch is as positive as the common, everyday light switch. If these cryotrons are wired together in ordinary circuits, capacitance and inductance will prevent their inherent speed from being used. We therefore plate the metal into intricate forms on glass, very thin, so that there is almost no capacity or induction. Using thin films of silicon monoxide as insulation we plate layer over layer until we have a complex assembly on a microscope slide, equivalent to a vacuum tube chassis in information-handling ability but so thin that a finger tip cannot feel its presence. This method of assembly may not only get us our 1000 megacycles, but get it for us cheap, for in mass production such techniques of assembly are much cheaper than the classical wiring and soldering.
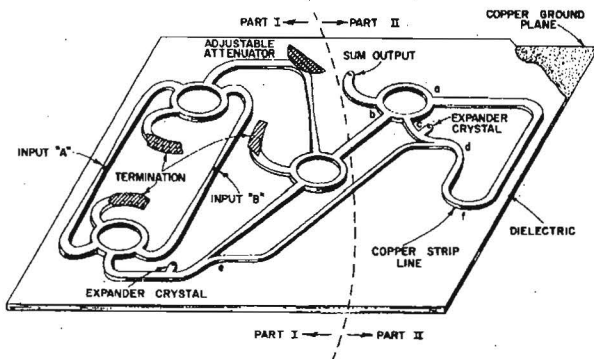

Fig. 1—A Full Adder

We also learned that when a complex aggregate of such units is assembled and inserted into its favorite habitat of liquid helium, which is kept in a double-walled flask with vacuum between its silvered walls, which in turn is kept in liquid nitrogen which is in another double-walled, silvered, evacuated flask, it had better be right. For if it doesn't work what can be done? You can't get your hands through the two-inch neck of the flask into the liquid helium, and if the mechanism is taken out of the helium it certainly won't work. Much experimentation is still needed.

We also learned that magnetic cores, an old stand-by, can be used in another way. By reducing the cores to flat thin films, much like the cryotrons, the "reluctance" of the core can be diminished and its speed increased. We already know how to make information machines from cores; BOGART is made mostly of cores. The thin magnets can be plated by the hundreds on microscope slides, a method much, much cheaper than threading cores, and the resulting information machine is faster than BOGART by about a thousand. This is somewhat short of our goal of 1000 megacycles, but cores have the advantage over cryotrons of working at room temperature. Since assembly is cheap, we may compensate for the insufficient speed by multiplying components and working in parallel.

Finally, we found out that improvements expected soon in transistors may make our speeds possible. But since we already see three other ways to our goal, we are leaving that one to be explored by others.

Phase I showed us four possible ways of getting the speed we want. Three of these we are exploring further in Phase II, now going on. Not only shall we get our speed, but we may also get cheaper techniques for building large information machines.

One last comment on the speeds we seek: an increase in speed by a factor of 1000 is an increase in capacity by a factor of 1000. This capacity will cost much less than 1000 times as much. At the present state of the art more speed is the cheapest way to get more capacity.

A kilo-megacycle computer could do one man-year of clerical work each second. The equivalent by hand would be a team of more than seven million clerks. One way of testing an information machine is to step it through its paces slowly enough for a human to follow. To do this for one second's work of a LIGHTNING machine we should have started in 1910 in order to finish now. The two outstanding innovations of the fifth decade of this century were the jet plane and the computer. A jet plane can go one hundred times as fast as a man can run. A computer can go ten thousand times as fast as a man can compute. LIGHTNING will go ten million times as fast.