June 13, 1956

COMPANY CONFIDENTIAL

THE MOTIVATION FOR THE HIGH-SPEED COMPUTER

John E. Griffith

In order to justify the large array of equipment which goes to make up the STRETCH computer, it is necessary to examine the problems that cannot be solved on present day machines. Now that 704's are being used, many customers are describing problems which the 704 cannot solve. Most of these problems are larger versions of problems presently on the 704, but some are actually new problems. For instance, several customers would like to invert 5000th order matrices. Others wish to reduce data at a rate of 100,000,000 words/day.

One customer regards the 704 as equivalent to 100 engineers for a certain problem in aircraft design, and wants a machine "equivalent" to 1000 engineers or more. This implies that computers can be used to amplify the output of a group of engineers, thereby effectively increasing the available manpower - an important point in the present manpower shortage.

The 704 is proving to be an extremely versatile machine for scientific and engineering computations, and it has generated a need for a much more powerful computer. It is therefore proper to examine a few of these problems to see what we want to do with a computer with the capabilities of the STRETCH machine.

As a background, it may be well to state that there are many problems that can be solved on computers, but not on present day machines. It may also be said that there is, at present, no hope of solving these problems except on computers.

Do we want to solve these problems? The answer is yes, and in some cases one has no choice. In the aircraft industry, the solution of a problem may decide whether or not a competitor's proposal can be overcome. If your competitor is using computers to solve problems, you have no choice but to use them yourself. In the AEC, the choice is predetermined by factors beyond the control of anyone in this country.

Do we have to solve these problems? The answer is yes again, but the reason is different. In the attempt to push forward into the future, computers allow us to attack certain areas that have resisted all other techniques. In certain cases, it is clear that we cannot afford to wait for other solutions to appear (Einstein is dead.) No matter how distasteful the computer approach will be, we have no choice but to proceed, if we are to progress as fast as possible.

Why cannot these problems be solved on the 704? Too little speed and memory are available. More than higher speed and larger memories are required, but these two parameters are the most important today.

Now let us look at a few of these problems in order to understand what we would like to do. I will take some examples from the thick forest of nonlinear partial differential equations. There are other equally dark and forbidding mathematical domains with many unsolved problems, but the one mentioned is under sustained attack by many people outside the AEC as well as within it, and will provide us with much new knowledge if we are successful.

These examples represent an almost fanatic desire to solve problems which are horribly difficult. This attitude is justified because in many cases the answer, when achieved, has proved to be worth much more than the money spent on the solution. The net return on solutions of some of these problems is often so great as to be fantastic. The intangible effects may be worth millions of dollars and in a few cases they have been estimated to prove it.

The first class of problem is one which occurs in the design of nuclear reactors. This class is interesting because it is solved by a technique which was invented especially for use on computers. It is called the "Monte Carlo" method and was invented by Stan Ulam and John Von Neumann. Basically, the technique amounts to throwing up one's hands and constructing a model of the system to be studied. But in this case, the model is statistical and is capable of an almost infinite variety of adjustment. (Slide 1)

A typical problem in this class is a 3 dimensional solution to a certain problem in neutronics. The problem involves 1,000,000 particles which must be followed in time for, let us say, 1000 time steps. Each particle will require 10 words of data for its description, thus occupying 10,000,000 words of storage. The computation time on STRETCH is 70 usec for each particle, all of which totals to 20 hours of machine time per problem.

This is a feasible problem, and easily worth the cost, for to get the answer by the experimental method (this particular problem may be solved experimentally in certain cases) would have cost at least one million dollars, by any estimate. Many more interesting cases cannot be solved experimentally or by any other method, and thus the answers constitute new information unobtainable without a computer.

This problem is also peculiar in that ten million words of random access storage is not actually needed. It is often more convenient to use tape. But the tape must be fast: on the average, one word every seven microseconds is needed for the full 20 hours. (Slide 2)

Another interesting class of problems is represented by hydrodynamics. The solution of wave motion, including shock waves, in three dimensions and time is a problem occurring in many places such as aircraft design, astrophysics, hydraulics, radio propogation, and solid state physics. A sample problem of this class might be one requiring (for 1% accuracy) a 100 x 100 x 100 three dimensional mesh. This totals one million points in space. One point requires 10 words of data to represent it, and thus, ten million words of storage is needed. The calculation time per point is 100 microseconds and for a problem of 1000 time steps, a total of 32 hours of machine time will be needed.

This, too, is a feasible problem, but it gives one a slightly uneasy feeling. However, this problem uses a technique with intriguing possibilities: weather prediction. In recent years, the average East Coast hurricane causes about one billion dollars damage. If hurricanes could be predicted sooner, a considerable saving in damages could be effected. It is interesting, therefore, to note that the problem described above could be designed to predict the path of a hurricane continuously during its life. Even if the machine predicts correctly only one out of ten hurricanes, it will save the taxpayer much more than the cost of the machine. Actually, methods presently under development promise to make the prediction of day-to-day weather a very practical possibility by the time the first STRETCH machine is produced.

There is a snag in the memory requirements for hydrodynamic problems: some problems can be done with tapes and a relatively small memory, but the more interesting problems are beginning to look like they will have the terrifying requirement of ten million words of high speed random access storage. (Slide 3)

The third class of problem is one of the worst that is known presently, but is seen more and more frequently, and seems to be on its way to becoming a permanent fixture in computing installations. This is the so called "composite" class of problems which is really a combination of the first two classes of problems.

As an example, let us suppose that we would like to solve a problem involving neutronics, hydrodynamics, electrical phenomena and mechanical motion over the same three dimensional mesh as used in the hydrodynamic problem discussed previously. The data to represent a point will take at least 100 words, thus requiring 100 million words to store the mesh. The calculation time for one point is about 3 milliseconds, and the computer time for 1000 time steps is about 7 months at 24 hours per day, seven days a week. This problem requires a total of 15 trillion arithmetic operations - a staggering total even for the STRETCH computer.

The fact that this problem is not feasible does not stop people from computing problems like these; they merely reduce the size of the problem until it is a feasible problem. At present, these problems are solved in one dimension with about 50 points. Needless to say, the largest errors are those caused by not computing in three full dimensions. Nevertheless, the answers represent information that is unobtainable under any other circumstances - including experimental circumstances in many cases.

Thus, we see the development of a trend: to replace experimental techniques wherever possible by computer techniques which are more general and which give answers to problems outside the range of experimental possibility. The extension of this technique to the overall prediction of the outcome of a set of interacting systems, such as weather phenomena or the operation of a jet engine, is obvious. We have shown that even the STRETCH machine will not, with all of its sophistication, solve all of the problems of immediate interest. But it is certainly a step in the right direction, for none of the problems mentioned here are possible on present day equipment, and we know that we need the answers to them. The answers to these large problems represent strategic information in many cases. If your competitor knows the answer to a certain problem, and he manufactured it on a computer, one has no choice but to compute it himself. If the problem is crucial, it becomes more important to try to keep ahead of one's competitors. It is generally impossible for anything to compete with a high-speed computer unless it is another high-speed computer.

Thus, the coming battle of STRETCH machines will be interesting to watch.

In closing, I would like to point out that there will be a need for the very best machine that can be built - and we propose to make the STRETCH computer fit this specification.

HYDRODYNAMICS

100 X 100 X 100 MESH

DATA = 10 WORDS/POINT

STORAGE = 10,000,000 WORDS

COMPUTATION = 100 μSEC/POINT

1000 TIME STEPS = 32 HOURS

COMPOSITE

100 X 100 X 100 MESH

DATA = 100 WORDS/POINT

STORAGE = 100,000,000 WORDS

COMPUTATION = 3 MILLISEC/POINT

1000 TIME STEPS = 5000 HOURS

$\approx$ 7 MONTHS

MONTE CARLO

1,000,000 PARTICLES

DATA = 10 WORDS/PARTICLE

STORAGE = 10,000,000 WORDS

COMPUTATION = 70 μSEC/PARTICLE

1000 TIME STEPS = 20 HOURS