The LIGHTNING Program

BY R. L. WIGINGTON

~~Confidential~~

Miscellaneous subjects: manpower and technical management, technical generalizations, interaction with other research programs, applications, and goals are examined to note general accomplishments, shortcomings, and lessons learned. The importance of providing sufficient technical leadership to such a diverse program is noted. A few independent technical generalizations, applicable to any techniques, are discussed. Other miscellaneous observations concerning the conduct of the program are made.

INTRODUCTION

At this point in time, looking back over most of the LIGHTNING program and forward a short time to its conclusion, there are a number of general comments about it—both technical and managerial—both good and bad—which should be made. The participants in the program—both in NSA and outside NSA—feel that the money and effort were a good investment; they feel that worthwhile things have been accomplished and that a basis has been established for further progress. These same people also know that everything did not go perfectly and that in such a speculative endeavor as this has been the way to better results is always more clearly seen after-the-fact.

INTERNAL MANPOWER AND METHOD OF TECHNICAL MANAGEMENT

Since the program began in the first phase without a desire for detailed technical direction by NSA, little manpower or time went into managing it during the first eighteen months. Considerable effort by the senior people of R/D had been necessary to formulate and launch the program, but they could thereafter monitor the program by semiannual formal reviews. After assignment of technical specializations to the contractors at the beginning of Phase II, NSA engineers in those specializations were assigned to keep informed of the contractors' technical activity, progress, and problems.

At this time the involvement of NSA began to become more direct. Internal research activities began to be arranged so that they were complementary to the contractor programs, and they began to form a broad base of technical capability which could be called on when needed.

As the program progressed into greater detail on fewer technical specialties in Phase III, and as the pressure built up to achieve a firm demonstration of research feasibility of the ideas under study, the NSA technical representatives further increased their frequency of contact with the contractors and their efforts in detailed technical direction.

Throughout the program the amount of NSA technical manpower allotted directly to the program was very small. At no time were there more than a few persons directly involved, and some of them were not on the project full time. These NSA technical representatives could not have coped with their responsibilities without the support of the internal NSA technical capability that has been built up, partly independently and partly as a result of the stimulation of the LIGHTNING program. The value of having this reserve of general technical capability in the field, to be applied as needed, was clearly evident throughout the program.

A greater commitment of NSA technical manpower earlier, in Phase II rather than in Phase III, might have resulted in quicker and more accurate recognition of how to confine the program to more details on fewer subjects and thus have permitted greater overall accomplishment. However, to carry this back farther—say into Phase I—would have been undesirable because it would have tended to destroy the fresh thinking demanded of the contractors.

As NSA moves to create capability based on this research, an increasingly direct involvement of NSA engineers will become necessary to guide the activity of contractors and to insure that this development activity is oriented as strongly as possible to NSA needs. So far, no one company has exhibited the best capability in all the fields —devices, logic circuits, memory techniques and assemblies, design and analysis—that are essential for building large scale equipment. Upon NSA scientists and engineers must fall the burden of being able to choose which set of capabilities best suits each application and of being able to exert technical leadership in each.

The original intention in the program was to narrow down the field—approximately half way through—to one contractor only. This was not done. At that time not enough had been accomplished, in detail, to make such a decision possible.

A fallacy, which grew into a major annoyance, was the phrase "*The LIGHTNING machine*", implying that the primary goal of the program was to be a single all-powerful ultra-speed machine. This program was too expensive for NSA if that were to be the end result. Of far more value is a wide-spread capability to be used in machines wherever high speed means something in terms of problem solution time. In this light, perhaps having more than one source for such techniques is a fortunate occurrence from the technical point of view.

The very nature of the program itself and its method of initiation were prime factors in how it was able to be conducted. First, it was a clearly labeled research gamble. Both the stakes and the potential gains were high. The best efforts of scientists and managers were expected to be able to go, as far as was possible, toward the very generally stated goals. Second, the impetus and early drive for the program came from the top down. It is doubtful that the necessary unity of purpose and momentum of action that were required could be generated by a working-level origin. Third, the program was conceived and conducted on a long term basis. On one hand, its top-level management support was continuous; no arbitrary cuts in the total resources available to the program were made; on the other hand these resources were regarded as a limit by those who conducted the program, who required that, within these limits, whatever research advances were made, should be backed up by sufficiently feasibile demonstrations that there was a firm starting point for the next step.

TECHNICAL GENERALIZATIONS

This program was conducted on the basis of the various electronic techniques which have been the focus of attention by different participants in the program at different times. Most of the budget requests, technical reports, and progress summaries have been outlined by such words as superconductivity, microwaves, magnetic films, and tunnel diodes. Yet the subject of this program is really fast computing or information processing and the electronic techniques, in general, useful for such purposes.

Some of the device-independent-properties of high speed equipment may not be so obvious when the concentration is on the amplifying element as we have usually presented the LIGHTNING story. Signals must be propagated from one place to another in the equipment, and the transmission of these signals is independent of the amplifying technique.

Transmission line engineering is not new to communication system designers and to radar and other microwave techniques people, but it has been a hard lesson for computer engineers to learn. They have now learned it and are prepared to use it in their design of fast equipment.

Another important general consideration is the matching of the amplifying mechanism to the transmission structure. One reason that cryotrons did not attain the speeds desired was that the switching device did not match the interconnection wiring in electrical impedance. This brought in new speed-limiting factors in addition to ordinary propagation delay.

We have always been acutely aware of propagation delay itself. Perhaps the most fundamental limit on machine speed is the speed of light. It sets the minimum time required to transfer information from one place to another. The speed of light in vacuum is 11.8″/ns, and in any real electrical circuit the propagation rate of information may be typically ½ to ⅔ of this. At every level of the hierarchy of present day computers there exist logical loops. Because of the limit on information transfer around such loops, their rate of operation is limited by the physical length of the path of the loop. Clever logical design can alleviate some of these problems, but an organization to eliminate all of them has not yet been found. The effect on logic design is clearly indicated in a paper by Hwang, Piel, and Raillard who state: "A logical decision made in one nanosecond and transported one foot is equatable to a logical decision made in two nanoseconds and used on location. This does *not* mean that ultra high speed computers must be built in a shoe box but rather suggests that logical designers should use space-time relationships as well as Boolean functions in the design of complex machines."[1]

For a long time we were misled by the "shoe box"theory. While we are not going to exceed the speed of light for signal propagation, this does not mean that miniaturization is an absolute necessity for an entire machine—only for the fast parts of it. Also, we have found that a machine designer can keep these additional restrictions in mind even though they do place a burden on him.

Looking in more detail at some of the techinque-independent speed-limiting factors, we observe that, in addition to the basic time constant of a phenomenon and the delay due to the speed of propagation, further delay is caused by stray reactance introduced in construction of devices and circuits and impedance mismatches in transmission paths. In order to make use of a phenomenon, it is necessary to provide electrical, magnetic or electromagnetic coupling to a material which exhibits this desired phenomenon. Known methods of coupling involve physical structures which have inductance, capacitance and resistance "strays." Until recently, the effect of many strays could be largely neglected bebecause their effects were generally small campared to the time constants of the phenomenon involved, or they were masked by the impedance level of the device. At 1 gc a 1 nanohenry inductance associated with a lead-in wire has an impedance of 6.28Ω. This impedance is not significant compared to the output resistance of a transistor which may be one thousand ohms, or to the output capacitative reactance due to the junction capacitances of the transistors. On the other hand, an inductance of 1 nanohenry in the lead to a tunneling junction with an internal resistance of ½ ohm and a capacitance of 10 pf is the major speed-limiting factor.

Whether the inductance strays or the capacitance strays are the speed-limiting factors depends on the ratio of inductance to capacitance, the internal resistance, and the impedance of the environment. If the reactive elements are in the proper ratio to match the resistive elements, their detrimental effects on speed are minimized. Unbalance results in either RC or L/R time constant limitation, depending on which of these time constants is greater. Since the speed of operation of a device generally increases as gain required decreases, it is desirable to minimize the gain required. This can be done by matching internal and external impedances to provide maximum power transfer.

Although transmission lines have always existed in digital equipment, their distributed properties have never before been so important. Design and analysis based on lumped element concepts lose their validity when the spectrum of wave lengths involved in the signals in the machine extends to dimensions large enough to be comparable to the physical size of the structure involved. Since electronic speed of operation has gone up faster than the size of structure has decreased, the distributed nature of real components and interconnections has become very important.

Because of the distributed nature of the interconnection paths, it is necessary to match the characteristic impedance of the transmission paths connecting elements as well as the elements themselves to obtain maximum power transfer. One additional effect of impedance mismatch is also important. Power which cannot be transferred into a load because of impedance mismatch to the transmission line is reflected back down the line toward the driving end. The operation of the "sending" element may be distributed by this reflection at some time after the original operation. Rereflection of the reflected signal of the "sending" end may cause further problems at the receiving end. If it is necessary to wait for such multiple reflections to die out, time is wasted.

Undesired coupling between points can occur through other stray reactances. Common path mutual impedance between points can also result in undesired cross-talk. Both of these difficulties are increased as frequency increases. These cross-talk problems are not new in character—only greatly increased in extent.

---

[1] In a paper prepared for the IFIPS Conference in Munich, 1962.

As a consequence, transmission line engineering must permeate all phases of the electrical analysis and design of ultra high speed computers, from power supply distribution to signal path routing. For example, very low (.05 to 2.0 ohms) characteristic impedance power supply distribution lines are needed to prevent stray signals from building up and causing undesired coupling. Since power supply voltages are low (90 mv to a few volts) and currents may be high (many amperes) in an equipment of any size, the D. C. resistance must also be low. Three types of signal propagation lines are in use. Various versions of multi-layer strip lines, miniature coaxial cables (O.D. less than 0.020 inches), and wire over ground plane structures are satisfactory. The impedance levels used range from 10 ohms to 150 ohms.

At times, brute force is used to get high speed; i.e., if there are stray capacitances to charge, a large current will be provided to do this quickly. However, the limit of heat removal from the assemblies and the temperature limits of operation of materials and phenomena prevent this from being extended indefinitely.

IMPACT ON AND ASSISTS FROM GENERAL HIGH SPEED COMPUTER FIELD

As is noted in another section of this report (Section IX—Scientific Communication and Impact), the scientific publications emanating from this program represent the major block of publication in the ultra high speed digital electronics field in the last few years. As measured by this publication activity and direct observation of the discussions which it generated at various technical conferences, the objective of the program to stimulate industry in this field has been very successful. Also, as noted in Section IX, the program received a copious return in the form of technical information and developments reported in open literature. Without these additions from the outside, the program would not have succeeded as well as it did.

One may wonder if any real "break-throughs" were actually caused by the existence of the LIGHTNING program in so far as it did seize every opportunity to use fundamental electronic properties without regard to origin. In the outline below of the origin of many of the basic processes and devices emphasized during the program, it is evident that ultra high speed computing was indeed the main subject of the program, as was intended, and not, as the casual observer may have thought, the discovery and study of new basic phenomena.

The usefulness of thin magnetic films for computer applications was recognized long before the LIGHTNING program started. NSA supported work on the subject at NOL Corona, California, and at Remington Rand Univac for several years before LIGHTNING added the incentive to look for really fast operating speeds. The cryotron was invented by Dudley Buck, and NSA had a cryotron research project which was being conducted before LIGHTNING. Again, the element that was added was a quest for speed plus a big push for developing the ideas to a state of usefulness. Microwave techniques of various types for computing were suggested before LIGHTNING started, and NSA was also in the middle of developing a small effort in that field when the LIGHTNING program was launched. Therefore, the notion that magnetic films, superconductive phenomena, and microwave techniques were of potential usefulness was not new to the general computer research field nor to NSA.

When Esaki announced the invention of the tunnel diode in January 1958, it was recognized as a potentially useful high speed component, and much work was done at many laboratories to investigate its properties and uses. Although the major emphasis of only one laboratory in the program, RCA, was directed to tunnel diodes and their application, a little work was also done at Sperry-Rand and IBM. The LIGHTNING contractors were in a position to recognize and make use of their properties. Of course, outside the program, a large amount of work was done at several laboratories and some work at many others.

Upon entering Phase II, some decisions were made concerning transistors and transistor circuitry. As noted in Section III of the overall report (Phase I), all contractors identified transistor techniques as important for high speed computing. Noting that other interests (principally NIKE-ZEUS) were already exploiting this field sufficiently, the decision to depend on those interests was a specific indication of a policy to profit by "outside" work wherever possible.

Scientific publication, therefore, was a two-way street so far as the LIGHTNING program was concerned. On the one hand, by contributing the major block of publication in the ultra high speed digital field, it showed others that very high speed is achievable, contributed a large number of scientific advances to the field, and encouraged others to do research in the field. On the other hand, it received in return very important assistance from work done outside the program.

APPLICATIONS

Since, as a research program, LIGHTNING is just ending, its full-scale pay-off is still some time away. The REDMAN program is viewed as the backbone of NSA's high speed development activities. Two studies have already been completed in R53 in the early build-up of activity on that program. For details and results, the reader should refer to the reports referenced below; the abstracts of those reports are quoted here to give an indication of their contents. One report is an evaluation of several sets of high speed logic circuits by the design of a particular functional assembly with all types. The summary of this report is:[2]

"Present attempts to develop an ultra high speed data processor using modular building blocks utilize a wide spectrum of circuit configurations and packaging techniques. This paper presents a start in the attempt to characterize the system limitations resulting from the characteristics of the individual modular building blocks and the method of proposed interconnection."

"The evaluation is based on logical and packaging considerations and does not consider circuit specifics such as component tolerance, power supply requirements, etc. A six bit variable mod accumulator was chosen as a small system test vehicle. Serial mode processing is not considered as it does not give maximal processing speed at the system level. The paper contains discussions and conclusions, where valid, on the relative merits of techniques used in the families of modular building blocks studied."

The other report looks at certain operational equipment and extracts the logic circuit usage and requirements imposed by the particular equipment. The abstract of this report is:[3]

"The basic purpose of this report is to present Special Purpose equipment logic requirements as determined by a study conducted on Special Purpose equipment in MPRO. These requirements are principally in terms of fan-in and fan-out. It is concluded that an "and" gate fan-in of three and an "or" gate fan-in of four with a minimum fan-out of five is the practical minimum requirement. Two other phases of the report are in chart form and are intended to display special purpose physical characteristics and requirements. One chart is an equipment construction time chart and the other an equipment facts chart. A final phase of the report contains the consensus of a discussion with C4 personnel on Special Purpose Systems Requirements. Feelings concerning the significance of speed, reliability and cost are discussed. These feelings tend to indicate a strong desire for high reliability with due consideration for high speed and reasonable cost."

---

[2] Donald J. Chesarek, "Logic Evaluation of Some Prepared Ultra-High Speed Modular Building Blocks," *Technical Memorandum R53/29*, Unclassified, June 1962.

[3] James J. Mihalik, "A Logic Evaluation of Analytic Equipments," *Technical Memorandum R43/30*, Confidential, June 1962.

These two reports are typical of what would have been extremely helpful earlier in the program in identifying worthwhile hardware achievements.

Some small-scale design studies were initiated at both RCA and Sperry-Rand Univac in September 1961 to see if the techniques to be demonstrated in the feasibility test systems would be applicable to real problems. Problems were proposed to the contractors in the areas of data comparison and counting and in summation of products calculations (such as are important in correlation and power density calculation). The aim was to test designability rather than to end up with an operational equipment. These studies will be published in another section of this report (Section VIII) upon completion. Trial designs are very important background work for moving on to the design of operational equipment. Too often, techniques are "accepted" or "rejected" on the basis of preconceived notions or shaky extrapolations from old applications. LIGHTNING was occasionally a victim of this failing. Both the novel difficulties presented to the designer by new hardware techniques and the increased opportunities for new design tricks which are also made possible by them are not obvious until several designs based on them have been completed. Even if these are only skeletonized paper designs, the tackling of several typical examples is worth the effort from the standpoint of accuracy of evaluation and confidence in predictions. Of course, such studies must be coupled with the actual construction of something that experimentally verifies the design assumption made.

This type of investigation was undertaken on LIGHTNING and on REDMAN—but much later than it should have been. This is a particular area in which more technical effort, within NSA earlier in the program, would have helped to achieve more pertinent research progress.

There have already been, and soon there will be more, pay-offs to which the LIGHTNING research has contributed.

A problem developed in the address selection matrix in the high-speed memories in Harvest; specifically, the speed of recovery of the diodes was not fast enough. A solution was found in diodes (1N921,22) manufactured by the Sperry-Rand Norwalk Semiconductor device plant. These diodes are silicon high conductance avalanche diodes not available from other manufacturers. Although an unbroken path of the development for these particular diodes has not been traced, it is highly probable that this is one of the unplanned and unknown NSA pay-offs from early LIGHTNING research. It may be recalled that, during Phase I, research at Sperry-Rand Norwalk provided the first "LIGHTNING speed" component, an aluminum-bonded silicon avalanche diode which would switch faster than could be measured at that time (0.3ns).

A design is now being studied at Sperry-Rand Univac for a special purpose machine called THORNHILL. Some of the same people are involved or have been involved in the design and construction of the magnetic film LTM. Many of the construction techniques developed for LTM, particularly the multilayer strip line back-panel wiring, are expected to be used. The technology established on LIGHTNING and previous NSA research on magnetic films led to the BARCROFT study and test memory. That basic design will also be incorporated into THORNHILL.

Uses of nanosecond techniques will be discovered which will benefit all areas of electronic data processing. Commercial emphasis, in the near future, will be on the ability of nanosecond techniques to provide a greater number of elementary operations per dollar per unit of time.

These techniques will probably first be widely used in commercial machines, constructed primarily with low speed hardware, to reduce mismatches of operation both within and between levels of the computer hierarchy. Such things as fast carry networks, special comparison networks, high speed scratch pad memories, and high speed multipliers will be designed for computers with microsecond memory cycles. Such uses avoid some of the problems of assembling large numbers of nanosecond circuits. Although nanosecond components are expensive at the present time, careful selection of applications for them will provide much improved performance at little increase in total cost.

The first commercial computers which can really be called "Nanosecond" computers (using nanosecond techniques except where speed is limited by connections to the outside world) will probably be serial computers. Most large general purpose computers are largely parallel machines, operating on the n bits of a word in a single memory cycle time of from 1 to several microseconds. Nanosecond techniques can be used in a serial organization to accomplish the same effective overall operational speed as the conventional "microsecond" machines, with considerable saving in the amount of hardware.

This commercial activity will complement NSA's development efforts and assist it in getting costs down. We have continually emphasized to the contractors that we cannot afford to be the sole customer for a given technology.

Machine organizers and equipment designers are presented with a challenge to make the most effective use of ultra high speed digital techniques, as represented by the accomplishments of LIGHTNING and related research activities in the same field. Some attention was given, during the program, to how best to use them. However, this part of the research was not developed to the extent that the hardware technology was. There is much unfinished business there. The feasibility test equipments, for the most part, are very closely related to standard machine organization and design from the mathematical point of view. The most novel thinking was done in connection with the research on superconductivity at IBM. The basic element in that case was so inherently different that a different organization became necessary. Unfortunately, in that case the technology didn't come through —or at least hasn't yet.

Which should come first, the specific machine requirements or the stock of electronic techiniques by which to satisfy them? As with all "chicken and egg" situations this can lead to endless arguments. The approach that has been taken in this case, good or bad, is this: based on a general awareness of the information processing requirements of NSA, an ultra high speed capability was demanded. The electronic techniques problem was attacked in greatest detail first, and a "research feasibility" has now been established. There remain the problems of translating this feasibility into an actual hardware capability and of increasing the machine organizational effectiveness to make the best use of that capability.

GOALS

Let us consider briefly the very pointed question of "how near have we come to meeting the LIGHTNING goals." The first difficulty is in reducing the original statement, "1000 Mc computation," to specific numbers with which we can compare.

Although they are not meaningful, in that no one would build such a machine, the most concrete numbers come from a simple linear extrapolation by a factor of one thousand of the IBM 704 or Univac 1103A, machines of five or six years ago. These machines, called megacycle machines, involved electrical rise times of 0.2–0.3 μs, shift times of 1 μs per bit, and memory cycle times of 10–20 μs. The 1000X scaling gives 0.2–0.3 ns, 1 ns, and 10–20 ns respectively. Electrical rise times of 0.2–0.3 ns have been achieved; shift times of 2 to 5 ns per bit have been demonstrated; and memory cycles of 25 ns for a very simple memory to a sizable memory now operating at 170–200 ns (expected to come down to or below 100 ns in the next six months) have been demonstrated. As a gross generalization, we may claim that we are within a factor of from 2 to 5 from the goals as expressed on this basis, with *research feasibility* firmly established and with considerable insight into how to go farther. This is not bad when aiming for an improvement factor of 1000.

Is there any other basis for reducing to numbers the degree of improvement in the figures of merit of computers or information processors that we have made possible? Although the "megacycle" is not a proper unit with which to measure the speed of a computer, we do not know a really good one. Another good way to start an argument would be to ask for proposals for such a measure. Actually, such a universal figure of merit may not exist. Digital machines and systems of machines are really very complex—both electrically and in the manner of their use. Any measure of effectiveness is a very sensitive function of the problem, or set of problems, to which it is applied.

With respect, then, to the question of satisfying the goals, we must be content that we have come reasonably close, as defined by the very poor measure of effectiveness that we have to use. We hope that within the period of the LIGHTNING research program we have lived up to General Canine's demand for "substantial progress."