Reliability and Characteristics of the Illiac Electrostatic Memory*

J. M. WIER†

Introduction

THE ILLIAC MEMORY was completed in the spring of 1952. It was tested during the early part of the summer and in August was attached to the previously completed arithmetic unit. About three weeks were needed to complete the physical work necessary for this and to eliminate all unsoldered joints, misconnections, etc. Detailed records were kept after this time. This log begins on September 3 and covers every "on" period in detail. Later entries are more specific due to the added experience gained in isolating difficulties.

The memory is of the Williams type. It is a parallel device with 40 words of 1,024 bits each. The storage tube which has been used is the 3KP1. It is operated with 2,000 volts accelerating potential. The regeneration amplifier has a gain of about 70,000 to a positive pulse and is a very straightforward RC coupled type. The nominal bandwidth is 60 kc to 500 kc. The coupling time constants are 5 microseconds each. The period of a regeneration cycle, that is, the time to regenerate one 40-bit word is 19 μsec.

* This work was supported in part by the Office of Naval Research.

† University of Illinois, Urbana, Ill.

Philosophy of Operation

During the first year of its operation a number of changes and improvements were made. Since it was recognized at the outset that this would be done, it was felt desirable to set aside time for these changes as well as time for periodic checks of machine operation. The

8:00 AM to 12:00 noon interval was chosen for this. The running period is considered to be from 12 noon to 12 midnight, five days per week. The Illiac is also on from 8:00 AM to 12 noon and occasionally at other odd hours, but these times, although completely logged, are not considered as running time, even though some computational work may be done during them.

This report is limited to the running period. All time in this interval is accounted for including any unscheduled maintenance required in the interval. All known errors in this interval are recorded. For this paper an error is considered to be a known and detected failure. If two failures occur in five minutes, even though the second occurs while in the process of isolating the cause of the first, this will be treated as an additional error. Such errors make up the majority of short interval failures.

A number of tests are run at regular intervals to determine the operating condition of the machine. This is done about 5:00 to 5:30 during the running period. Special tests of the memory alone include a read-around ratio test for testing address interactions and flaw tests to determine the existence of nonstoring spots at addresses used. Since the Illiac never has possessed a complete set of flaw-free tubes, it has been necessary to operate on cathode-ray tubes which contain nonstoring spots. A relatively dc-stable address generator used in combination with small trimming magnets placed near the storage surface of tubes known to contain flaws has been partly successful in keeping storage points off of nonstoring spots on the cathode-ray-tube screen.

A further and more exhaustive test is performed by using a so-called leapfrog test.[1] This test is designed to detect errors of all types which may occur, with the exception of errors in the printing of results. Due to a closed loop through the output punch, even the input and output are checked. This code has the advantage of isolating memory errors from other types and giving their memory locations. It is through its use that this report is able to give the cause of a given error in most cases. There are, of course, errors that occur at other times which are not so well defined and these often are not specified as to cause.

One characteristic of the Illiac causes the existence of undetected errors to be much less likely than they might at first sight seem to be. Since the instruction code of the Illiac is not completely decoded, there are a number of possible arrangements of the instruction digits which are meaningless and upon whose appearance the machine will stop. Probably the most useful of these is the order which is represented by all zeros. Since the memory is cleared to all zeros before a problem is run, if something goes wrong with a code, it is usually shortly in one of two conditions, looping or stopped in the all zero region of the memory. It is then evident that such failures are easily detected and it often turns out to be a simple matter to locate the failure which has caused this behavior. Of course, random and very infrequent failures in the arithmetic operations are more difficult to isolate although it has been found in general that such errors soon make themselves more and more evident and do not go long undetected.

Since this report covers the memory it should be pointed out that memories in general have one property which makes the determination of an error quite simple. Unless they fail completely, they tend to remember their errors.

Collection of and Pertinent Comments about the Statistics

As has been mentioned this report covers the noon to midnight period for five days per week. Since the machine is turned off completely every night there is a larger probability of machine failure in the period 8:00 AM to 12:00 noon than in any other four-hour period. This is because of the disturbances caused by the turn-off-turn-on process. The error rate is in general largest then. This period is used to attempt to place the machine in such a state by noon as to run error-free the remainder of the time. The degree of success attained in this may be seen from the statistics which follow.

There are a few disturbing factors which make this report slightly pessimistic. Certain of the running periods were immediately after major engineering changes and required troubleshooting to iron out errors caused by such prosaic things as unsoldered connections or, occasionally, logical errors. There are three such periods rather carefully distinguished, the first one being in the first two weeks after the final assembly of the machine; the second one after the complete replacement of the adder by a more efficient type; and the last one by an operational change in the memory which improved the read-around-ratio figure. The adder replacement probably does not affect the memory figures as much as the others.

The figures given encompass only errors which were reasonably assignable to the memory. It is possible that a few others may have been due to the memory, since about 80 errors in this period could not be traced.

Statistics

These statistics cover a period starting on September 3, 1952 and ending August 8, 1953. During this time the Illiac logged 4,100 dc hours. This report covers 2,976 hours of this, falling in 12 noon to 12 midnight interval.

During this time a total of 479 known machine errors occurred. Of these 179 were directly assigned to the memory and 80 were not traced as to cause. Thus the memory caused a minimum of 37.4 per cent of the total number of errors and a maximum of 54 per cent of them.

The errors attributable to the memory were investigated and it was found that 123.2 hours of computing time were lost due to the memory. This lost time includes all the time lost due to running the problem

[1] D. J. Wheeler and J. E. Robertson, "Diagnostic programs for the Illiac," Proc. I.R.E. vol. 41, pp. 1320–1325; October, 1953.

terminated by the error, the time necessary to make repairs, and the time to make sure the trouble had been cured. The latter process usually consists of the running of test programs covering the problem at hand. This lost time is 4 per cent of the total period.

Since there were 179 memory errors in 2,976 hours and 80 errors not traced, mean free time between memory errors is 16.6 hours at best and 11.5 hours at worst.

Although mean free time figures are instructive, it was felt to be desirable to analyze the errors more completely. This was done by making a table of all known memory errors in this interval, noting the time of each error. Then the time between successive errors was found, considering each 12-hour period to be immediately succeeded by the next 12-hour period. These differences then indicated the different times between memory errors. These differences were grouped by length of error-free run and counted. Fig. 1 shows the results of that process. Note that the time intervals are growing exponentially larger to the right. All curves are normalized to the total number of memory errors. It should be noted that the 8–16 hour group is the largest.

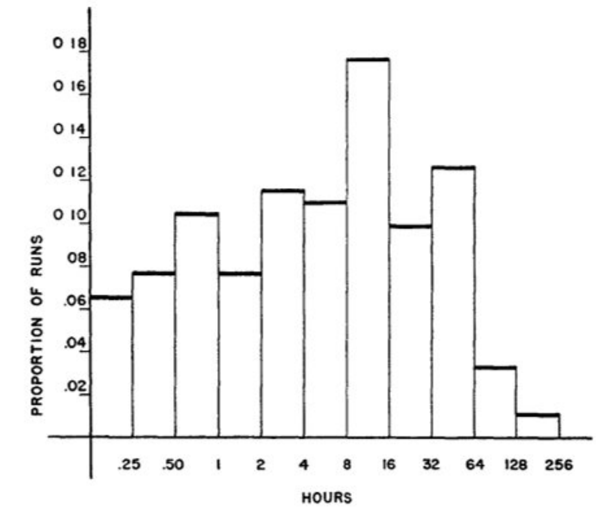

Fig. 1—Length of run distribution.

This is probably due to the disturbing influence of turning the machine off between each successive set of 12 hours. The error-free intervals from 0 to 1 hour are largely in the unscheduled maintenance periods.

Since the ordinates are normalized the height of any given group bar is the part of the total number of memory errors in each group. In order to find the probability that a given memory run will be longer than some time, we may add all the groups to the right of that abscissa. The results of doing this are shown in Fig. 2. Thus if one takes a time immediately after an error occurs and wishes to know the probability of running a problem of length H hours, this may be found by going to the abscissa H and reading off the probability on the axis of ordinates. It will be noted that the point at which the probability is $\frac{1}{2}$ is somewhere in the 4–8

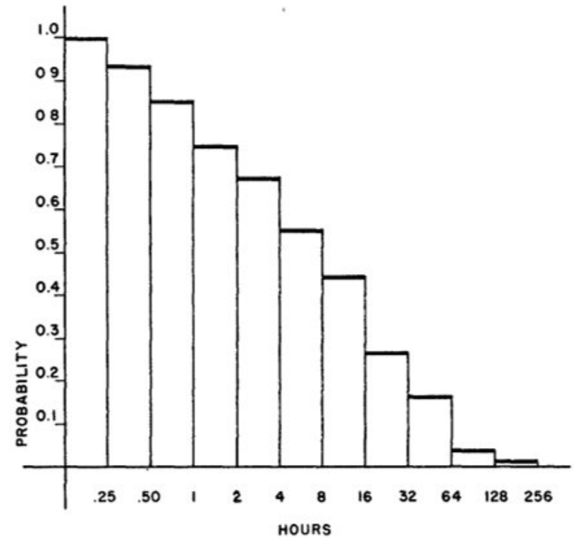

Fig. 2—Probability of a given run length.

period, being nearer the latter. The fact that this does not agree with the 11 to 16 hour mean free time stems from the large number of closely spaced errors in the unscheduled maintenance period.

Causes for These Errors

In order to determine what caused these failures, each error was catalogued as to a type of error. Table I shows how these errors are distributed:

TABLE I

Causes for Errors in the Memory

Cause	Number of Errors
Flaws	81
Unknown	71
CRT Intensity Drift	10
Poor CRT Focus	7
Address Generator	6
Changes Logic Circuits	2
Noisy Regeneration Amplifier	1
Poorer than Stated Read-Around Ratio	1

It will be noted that there is a large contingent of unknown errors. Most of these were accumulated in the early history of the machine when the operators were not so expert in finding or analyzing the errors. Also, the leapfrog test code in use then was not so completely able to isolate the trouble. It is evident that by far the commonest error involves flaws on the cathode-ray tube surface. As has been said, the Illiac operates in the presence of these flaws and from time to time components drift in the address generator causing the difficulty noted. Other troubles are few and are usually fairly simply repaired. It should be noted however, that the largest part of the trouble has been due to the cathode-ray tube, a component which has also had the most effort placed in the direction of its improvement by various manufacturers and computer groups. All of the tubes of an improved nature which were tested were found to be capable of eliminating the majority of these errors. These include sample lots of the IBM 79, the RCA C73376, and C73621 types.

After diagnosing these errors, various corrective actions were taken. The most frequent of these was to place a small magnet near the tube face to move the raster off the flaw. Although this is felt to be somewhat makeshift procedure, it has been reasonably successful in holding flaw difficulties down.

The original set of 40 cathode-ray tubes was selected from a set of 160 standard 3KP1's. About 50 per cent of this number were considered to be good enough for retention as possible storage tubes. The remainder were discarded either for the presence of more than two non-storing flaws or for reasons of poor read-around ratio. Forty of the better tubes were placed in the memory for preliminary tests before connection to the Illiac. The remainder were used as replacements. Twenty-five additional tubes have been procured since then for replacement purposes.

During the period September 3, 1952, through August 8, 1953, it was necessary to replace 43 cathode-ray tubes. Table II shows the distribution of difficulties for which these changes were made. It is evident that the principal reasons for removal were flaws and low read-around ratio; these accounted for 30 of the 43 replacements. The read-around ratio removals were made to improve that quantity. Tube changes were made for reason of flaws only when two nonstoring or marginal flaws existed at storage points, at the same time.

TABLE II

Cathode-Ray Tube Faults

Fault	Number of Removal Caused by Fault
Flaws	16
Rar	14
Suspicion	8
Poor Focus	1
Shorted Elements	1
Loose Aquadag Connection	1
Open Filament	1
Remote Cutoff	1

Of the remaining 13 tubes, 8 were replaced on suspicion that they were causing difficulty, but without very rigorous evidence. It is quite possible that some of these replacements were unnecessary.

The chassis housing the regeneration amplifier and circuitry necessary for the logical operations needed to write in and read out of each position was a relatively trouble-free device. As may be seen from Table III, it was found necessary to remove 16 chassis for repair in the September 3, 1952 through August 8, 1953 period. Those removed on suspicion were generally found to have poor amplifier tubes which resulted in an occasional noise spike coming out which resulted in an error. In this sense they are indistinguishable in symptoms from those labeled noisy amplifier and microphonic. Only one other failure was found in this category and that was a leaky bypass condenser.

Those errors listed in the logical section were errors which completely prevented storage. Four of them were caused by open filaments and one by a grid-to-plate short.

Thus the failures in the memory were almost exclusively tube failures.

TABLE III

Regeneration Chassis Faults

Reason for Removal	Number of Chassis out for this Reason
Suspicion	8
Logical Section	5
Microphonic	2
Noisy Amplifier	1

The Address Generator

The address generator forms the analog equivalent of any given digital address and deflects the cathode-ray-tube beam to that address. It caused six errors. Four of these were due to tubes with open filaments and the remainder were grid-cathode shorts.

In addition to those actions taken after failures had occurred, a considerable amount of preventive maintenance was done. This maintenance included tolerance tests of various supply voltages, hammer testing of circuits to locate microphonic tubes or circuits, and a complete test of all tubes except those in the regeneration chassis, at about 3,500 dc hours. This latter check resulted in the replacement of a relatively large proportion of the type 5687 tubes which are used extensively as drivers.

It should be mentioned that the set of ten pulsers and their output drivers caused no known trouble in this period although a number of tubes were replaced in them at the time of the wholesale tube check. Certain tubes in the address generator were also removed at this time without their having caused any known error.

Read-Around Ratio

One figure of merit for all Williams tube systems is the read-around ratio. Since there is some interaction between adjacent storage spots in the tube, too frequent consultation at some given address without regeneration of its neighbors tends to cause loss of information in these neighboring addresses. Of course, as the packing density increases, this interaction becomes greater. It appears that 1,024 storage spots per tube is the practical limit within a factor of two with present tubes used in a parallel manner.

In order to improve the read-around ratio figure, various efforts have been directed toward improving the cathode-ray tubes and the method of operation. The Illiac memory system, limited to standard cathode-ray tubes because of procurement difficulties, has used circuit improvement. R. Thorenson of SWAC proposed last February that better results might be achieved by sensing the presence of information at a later time in a typical output pulse.[2] A scheme using this information in a way slightly different from that proposed by Thorenson was devised for the Illiac. After tests in a test rack this was found to be sufficiently better than the normal Williams system to make installation in the Illiac worthwhile. As the changes were minor, this took only a weekend. The result was a net improvement by a factor of three to four in the read-around ratio.

The test for the read-around ratio is programmed and tests every point on the raster in the following way. The half of the raster to be tested is filled with 0's. 0's correspond to the longest beam on time with the Illiac system. The addresses surrounding the test address are filled first with 0's and then the test address is bombarded n times with 0's, (where n is the test read-around ratio number). The surrounding addresses are checked for failures and these are noted, if any. This is done twice at the address to avoid the possibility of stray regenerations making the test less severe. Then the surrounding addresses are filled in with 1's and the bombardment and checking takes place as before. Failures are printed out with suitable indicating information to give the test number n, the direction of failure, (from 0 to 1 or 1 to 0), and the address and tube failing. This test for failures both from 1's to 0's and 0's to 1's is unnecessary with the ordinary Williams system operation since failures occur only in one direction. With the Illiac system, a more balanced one, both failures occur with similar frequencies.[3] The point of marginal failure is 3 to 4 times better than the normal Williams system. This system has been used on the Illiac since March 1953 and since then the read-around ratio value from the standpoint of the programmer (at which no failures occur on any tube at any address) has varied from 60 to 80 depending on the condition of the tubes in use. This has been enough to free the coders of excessive care in coding against read-around ratio. It is possible to cause failures if sufficiently short, fast loops are used.

[2] R. Thorenson, "An Improved Cathode-Ray Tube Storage System," *National Bureau of Standards Report 2275*; February 6, 1953.
[3] J. M. Wier, "The Illiac Memory System," *Proc. of Symposium on Large Scale Digital Computing Mach.*, Argonne Nat'l. Lab.; August, 1953.

Summary

The following statistics are of greatest significance. Machine time lost due to memory failures totals 4 per cent of 2,976 hours of running time. A total of 179 errors were directly attributable to the memory for a mean free running time of about 16.5 hours. Starting after an arbitrary failure of the memory, there is a probability to $\frac{1}{2}$ that a continuous error-free run of at least 4 to 8 hours will follow, the true figure being somewhat closer to 8 hours. The read-around ratio is a minimum of 60 to 80 on all tubes.

In the September 3, 1952, to August 8, 1953 interval, including time not called running time, all cathode-ray tube replacements totaled 43. It was necessary to repair 16 regeneration amplifier chassis in this same time. Other difficulties were few and simply repaired.

Further Improvements in the System

As the principal cause of failures has been the cathode-ray tubes, a complete set of the RCA type C73621 has been ordered and partially received. In order to eliminate the other errors a regular system of checking vacuum tubes has been instituted to catch tube failures before they cause machine errors

A number of experiments using more elaborate discriminating methods have been tested to make the system more independent of operating variations. The most successful of these involves an integration sensing method which integrates the magnitude of difference between the desired signal and the signal present over the whole meaningful output pulse. This system is very insensitive to noise in the system but involves somewhat more circuitry than can be added to the present regeneration chassis.

It is felt that the system has reached a state of reliability compatible with the error rates of the rest of the machine with the exception of the state of the cathode-ray tube. Further improvements, while desirable, cannot add too effectively to the stability of the system as a whole, since the present Williams memory is nearly as reliable as the average for the remainder of the machine, judged on a per tube basis.

Discussion

L. E. Kanter (IBM Corporation): For every reference to the memory is there a fixed regeneration cycle period or how does the time allotted to execution of instructions compare with time allotted to regeneration by the ILLIAC?

Mr. Wier: The memory is basically driven by a clock which has a period of 19 microseconds. In a sense the arithmetic unit is a slave of the memory. The arithmetic unit waits until it reaches the time in the memory cycle at which it can acquire access to the memory. However, it is so arranged that it is impossible to require that the memory give forth two numbers in two successive regeneration periods. For this reason the upper value on the read-around-ratio to which it is possible to subject the machine is limited because every other period at the very least has to be a regeneration cycle.

J. L. Cochran (Defense Department): How is heater voltage applied to ILLIAC memory tubes? That is, the length of time from zero to full voltage?

Mr. Wier: During the period covered by this report, the voltage was applied abruptly. I will not defend this, however we have not had an excessive amount of trouble. Currently the voltage is applied as it has been all the rest of the time on the remainder of the machine. The voltage is applied through a large and nonlinear resistance which consists of a series of heater elements. The voltage starts out at essentially zero, builds up slowly to about 3 volts and then these heater elements are shorted out and the voltage goes abruptly to the value at which the filaments operate, namely, 6 v.

S. Greenwald (National Bureau of Standards): Would you give some detail on magnets used to avoid blemishes, such as shape, strength and location? What is the effect on CRT focus?

Mr. Wier: I am in a very bad position to answer this because the magnets which we are using are of the "friendly dog" type and I do not have detailed characteristics of these magnets. I think the only comment which might be made is that they are weak magnets and they are not standardized. As to their effect on focus we have found that there has been no lessening of any quantity which would normally be a function of focus such as the read-around-ratio unless the magnet is placed closer than about $\frac{3}{4}$ inch from the face of the tube.

Oliver Whitby (Stanford Research Institute): After the machine is first turned on in the morning, how long is it before the incidence of failures settles down to the average rate noted for afternoon operation?

Mr. Wier: The reason that I have avoided the mornings is because of the fact that we have made a large number of engineering changes and this process is still going on. This is not because the machine abruptly becomes a pile of junk at eight o'clock in the morning. When the machine is turned on, it may reasonably be expected to run any one of the test programs 5 minutes after it has been turned on. It usually takes this long to achieve a stable charge on the face of the cathode-ray tubes. I would say that probably the incidence of failure during this time is slightly higher than it would be during the afternoon, even after 5 minutes, but at the end of an hour the operation is substantially identical. The only reason that I have avoided the morning period is that I wished to avoid classifying periods which were not used for preventive maintenance but used only for changing the equipment.