Report 11

Quarterly Progress Report 11 | *Covering the Period 1 December 1962 to 28 February 19(*

GRAPHICAL DATA PROCESSING RESEARCH STUDY AND EXPERIMENTAL INVESTIGATION

Prepared for:

U.S. ARMY ELECTRONICS RESEARCH AND DEVELOPMENT LABORATORY
FORT MONMOUTH, NEW JERSEY

CONTRACT DA 36-039 SC-783
SCL 4087 (18 NOVEMBER 1958
FILE NO. 40001-PM-60-91-91(65(
DA PROJECT 3A99-22-001-02

By: *A. E. Brain* *G. E. Forsen* *D. J. Hall* *N. J. Nilsson*

STANFORD RESEARCH INSTITUTE

MENLO PARK, CALIFORNIA

*SRI

The enclosed report has been sent to you in accordance with instructions from the sponsoring agency. The report title, the author(s), the sponsoring agency, the contract number, and the SRI project number appear on the cover and title page of the report.

Editorial Office
Division of Engineering Research
Stanford Research Institute
Menlo Park, California

STANFORD RESEARCH INSTITUTE
MENLO PARK, CALIFORNIA

March 1963

Report 11

Quarterly Progress Report 11 | Covering the Period 1 December 1962 to 28 February 1963

GRAPHICAL DATA PROCESSING RESEARCH STUDY AND EXPERIMENTAL INVESTIGATION

Prepared for:

U.S. ARMY ELECTRONICS RESEARCH AND DEVELOPMENT LABORATORY
FORT MONMOUTH, NEW JERSEY

CONTRACT DA 36-039 SC-78343
SCL 4087 (18 NOVEMBER 1958)
FILE NO. 40001-PM-60-91-91(6500)
DA PROJECT 3A99-22-001-02

By: A. E. Brain G. E. Forsen D. J. Hall N. J. Nilsson

SRI Project No. 3192

Objective: To conduct a research study and experimental investigation of techniques and equipment characteristics of pattern recognition systems suitable for practical application to graphical data processing for military requirements.

Approved:

C. A. ROSEN, MANAGER APPLIED PHYSICS LABORATORY

J. D. NOE, DIRECTOR ENGINEERING SCIENCES DIVISION

Copy No. 54

ABSTRACT

A pattern may be presented for classification either as a code word of ones and zeros or of plus ones and minus ones; weight changing in the learning machine is effected by adding the pattern vector to the weight vector. It has been found, by digital computer simulation, that convergence is secured much more rapidly when the (+ 1, - 1) representation of the input patterns is used.

An intensive study has been made of the variability of the second-harmonic weights using a test procedure closely approximating the manner in which the cores will operate in the learning machine. Satisfactory performance was obtained over a useful range of drive levels.

A discussion is presented of the machine logic and timing, which have now been worked out in detail.

CONTENTS

ILLUSTRATIONS

TABLES

PURPOSE

It is the objective of this project to conduct a research study and experimental investigation of techniques and equipment characteristics suitable for practical application to non-alphanumeric graphical data processing for military requirements. All phases of the graphical data processing art will be considered, including the treatment of raw graphical data, identification, programming, selection, indexing, access to storage, and presentation. The studies and demonstrations of feasibility will be designed to evaluate the practicability of the proposed techniques and systems, with sufficient detail to be useful in establishing the design criteria necessary for equipment procurement.

The program of work to be carried out in accordance with the extension of Contract DA 36-039 SC-78343 will consist of:

(1) The study and development of organizations of combined fixed and adaptive networks that will permit recognition of patterns independent of size, displacement, and rotation, (a) in the presence of interfering signals and noise, and (b) on a real-time basis.

(2) The development of components and subsystems suitable for implementing the schemes devised in (1).

(3) The design and construction of an experimental Graphical Data Processing Machine making use of the techniques and components found to be most practicable by investigations (1) and (2).

PUBLICATIONS, LECTURES, REPORTS, CONFERENCES

On February 12 and 13, meetings were held at Stanford Research Institute, Menlo Park, California with the Sponsor's representative, William A. Huber of the Data Transducer Branch, Communications Department, U.S. Army Electronics Research and Development Laboratory, Fort Monmouth, for the purpose of reviewing experimental techniques and equipment for pattern recognition from graphical data.

I INTRODUCTION--SOME PRACTICAL PROBLEMS OF LEARNING MACHINE DESIGN AND CONSTRUCTION

MINOS II is now well advanced into the construction phase and all of the irrevocable major decisions have been made; fortunately, none of them have so far been invalidated by experience. It is inevitable that, when novel experimental equipments are being constructed for the first time, many decisions must be made on the basis of incomplete data--and the missing data does not exist anywhere. Also, one suddenly becomes aware that folklore has entered into the scheme of things and that one really does not know for sure that generally accepted relationships have a basis in fact.

For a long while, for example, we have believed that a (+ 1, - 1) scheme for changing weights leads to more rapid convergence than a (1, 0) scheme. Digital computer simulations have shown a trend in this direction over the years, but no proof has ever been offered, and there is at the back of one's mind the feeling that really they are more or less equivalent. Perhaps any difference is just an illusion. Which scheme should be used for MINOS II? There is a significant difference in the amount and complication of the hardware required to implement each scheme; the (1, 0) scheme is much simpler. Should we then go ahead and gamble on the (1, 0) scheme and perhaps miss a winner?

Fortunately, it was possible to resolve this dilemma by evidence obtained in an investigation being sponsored by Rome Air Development Center. The relevant results are briefly reported in Sec. II. The (+ 1, - 1) scheme is, in fact, very much better, and the additional expenditure in hardware is definitely warranted. It is perhaps salutory to recognize that we do not yet have a rigorous proof for the convergence of the majority logic training algorithm, whose efficacy is so well demonstrated in Figs. 2 through 7.

Another of the major unknowns, but one where direct practical steps may be taken to acquire the data, relates to the magnetic weights. How uniform are the cores in terms of our completely non-standard application? How critical is the drive current? Is it certain that an operating point exists that is valid for all of the cores in the matrix? The results of the experiments that were undertaken to provide answers to these questions are reported in Secs. III and IV.

Finally, the increment-decrement logic must be worked out in detail; when finally resolved, it is found to be appreciably more complicated than had been envisaged. Since this is to be an experimental machine, an attempt must be made to anticipate all the variants that people are likely to ask for in the years ahead--and to provide for them with no increase in cost. A brief review of the machine logic is presented in Sec. V.

II AN EXPERIMENTAL COMPARISON OF THREE LEARNING-MACHINE TRAINING RULES

A. BACKGROUND

In the design of MINOS II, it is necessary to specify whether the 100 binary inputs to the learning machine shall be represented by plus ones and minus ones (+ 1, - 1) or by ones and zeros (1, 0). This decision will affect the circuitry of the learning matrix; in particular, circuitry that can accommodate (+ 1, - 1) inputs is somewhat more complicated than that necessary to accommodate (1, 0) inputs (see Sec. V).

The standard error-correction training rule affects the input-output behavior of the learning machine in a significantly different way, depending on whether the rule is applied to a (+ 1, - 1) input machine or to a (1, 0) input machine. This difference is due to the fact that for a (1, 0) machine, only active weights (weights connected to a one input) are adapted, whereas for a (+ 1, - 1) machine, all of the weights are adapted. A limited amount of past experience has indicated that the (+ 1, - 1) machines converge faster than do (1, 0) machines. It was decided to conduct a series of computer-simulation experiments* to determine whether or not the training rate of (+ 1, - 1) machines was sufficiently faster to justify the added complexity of circuitry. In this section we shall describe these experiments and present a brief summary of the results.

---

*These experiments comprise a part of a larger study on Adaptive Mechanisms being performed under sponsorship of the Rome Air Development Center [Contract AF 30(602)-2943]. In view of the importance of this work to the design of MINOS II which is in its final stage, it was decided to report these empirical results in considerable detail here. When further work, including theoretical studies, is completed a full report will be issued for the Air Force.

B. DESCRIPTION OF THE EXPERIMENTS

To test the performance of the two methods, a one-bit output majority-rule learning machine was simulated on the IBM 7090 computer. This simulated learning machine was a scaled-down version of one of the six parallel units of MINOS II. A schematic diagram of the simulated machine is shown in Fig. 1. Twenty binary inputs are operated on by five threshold logic units (TLUs), which produce a one-bit output according to majority-rule logic. The five threshold logic units are connected to

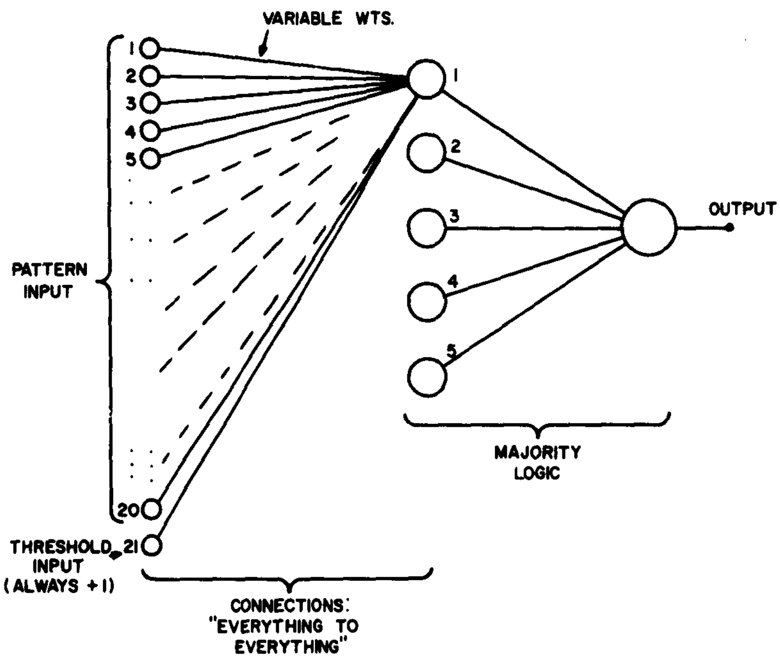

FIG. 1 SYSTEM ORGANIZATION FOR COMPARING (1, 0) AND (+1, −1) LEARNING MACHINES

the 20 inputs by an everything-to-everything scheme employing 100 adjustable weights. Each TLU also has an adjustable threshold simulated by adjustable weights connected to a 21st input, which always has the value (+ 1). The total number of adjustable weights is therefore equal to 105. The initial values of all weights and thresholds, before training, were in all cases equal to zero. [These initial conditions represent equivalent starting positions for both (1, 0) input machines and (+ 1, - 1) input machines.]

Learning curves were obtained for random sets of randomly categorized patterns, first represented by the (1, 0) scheme and then by the (+ 1, - 1) scheme. Six different random pattern sets of 90 patterns each were used. These sets were divided into three groups. In Group I, (Pattern Sets 1 and 2), each pattern had exactly five ones; in Group II (Pattern Sets 3 and 4), each pattern had exactly ten ones; in Group III (Pattern Sets 5 and 6), each pattern had exactly fifteen ones. Each group had two sets of different random patterns, and two learning curves were obtained for each set. One learning curve is the result of training a machine whose input patterns were presented as ones and zeros; the other curve is the result of training a machine on the same patterns presented as ones and minus ones. In addition, a modified training rule was used for the patterns in Group III. This modified rule attempted to train a (1, 0) input machine in such a way that its learning performance approximated more closely that of a (+ 1, - 1) input machine operating under the ordinary training rule. Therefore, for Pattern Sets 5 and 6, a total of three learning curves were obtained.

A total of six pattern sets were used for the following reasons:

(1) Three groups were chosen to see if the difference in training rates between (1, 0) and (+ 1, - 1) machines depended at all on the number of ones in the pattern.

(2) Two different sets were included in each group to give an indication of the differences in learning curves for different pattern sets with the same number of ones.

A total of 90 patterns were used in accordance with a (local) rule-of-thumb that the number of random patterns that a machine can learn in a number of iterations appropriate for a practical application is roughly equal to the number of adjustable weights per output bit. The simulated machine had 105 adjustable weights.

C. TRAINING RULES

The three training rules tested all had the following characteristics:

When the learning machine output is in error, a determination is made of how many TLUs must have their responses reversed so that the majority will vote correctly. Let the minimum number of such reversals necessary be equal to k. Of those TLUs voting incorrectly, one selects the k whose analog sums are closest to threshold and prepares to reverse their responses.

Suppose the response of the ith TLU is to be reversed: Then, its weights must be adapted. The three training rules differ in the way in which this reversal is accomplished:

(1) (1, 0) Training Rule [for (1, 0) input machines]--An increment is added to each active weight (a weight connected to a one input). The size and direction of each increment are the same for each active weight and are determined by the total change needed in the analog sum to effect a reversal of the TLU binary output.

(2) (+ 1, - 1) Training Rule [for (+ 1, - 1) input machines]--An increment is added to all weights. Those weights connected to plus one inputs are altered in a direction opposite to that of weights connected to minus one inputs. The size of the increments is the same for all weights and the size and direction is determined by the total change needed in the analog sum to effect a revers. of the TLU binary output.

(3) Modified (1, 0) Training Rule [for (1, 0) input machines]--An increment is added to all weights. Those weights connected to plus one inputs are altered in a direction opposite to that of weights connected to zero inputs. The size of the increment is the same for all weights and the size and direction is determined by the total change needed in the analog sum to effect the reversal of the TLU binary output.

The (1, 0) and (+ 1, - 1) training rules were applied to all six pattern sets, whereas the modified (1, 0) training rule was applied only to the Pattern Sets 5 and 6.

D. RESULTS OF EXPERIMENTS

The learning curves for each of the six sets of patterns are illustrated in Figs. 2 through 7. Each learning curve depicts the number of errors made (out of 90 patterns) during a test procedure conducted after each iteration through the pattern set. The following conclusions seem warranted as a result of comparing the (1, 0) rule curves with the (+ 1, - 1) rule curves:

(1) In all cases, the (+ 1, - 1) rule converges to zero errors faster and more directly than does the (1, 0) training rule.

(2) The disparity between convergence times for the (1, 0) and (+ 1, - 1) training rules increases with the percentage of ones in the patterns, being least noticeable for the case of 25% ones and increasing to a large factor in the case of 75% ones.

(3) The convergence time for the (+ 1, - 1) training rule is little affected by the number of ones in the patterns.

It can be shown theoretically that the modified (1, 0) rule would exhibit a learning curve almost identical with that of the (+ 1, - 1) rule when the percentage of ones in each pattern is equal to 50%. For this reason the modified (1, 0) rule was not tried on Pattern Sets 4 and 5.

Examination of Figs. 6 and 7 indicate that the modified (1, 0) training rule results in a learning curve whose convergence time is intermediate between those of the (1, 0) and (+ 1, - 1) rules. For this reason, the modified (1, 0) rule was not tested on Pattern Sets 1 and 2, where the (1, 0) and (+ 1, - 1) rules produced very similar curves.

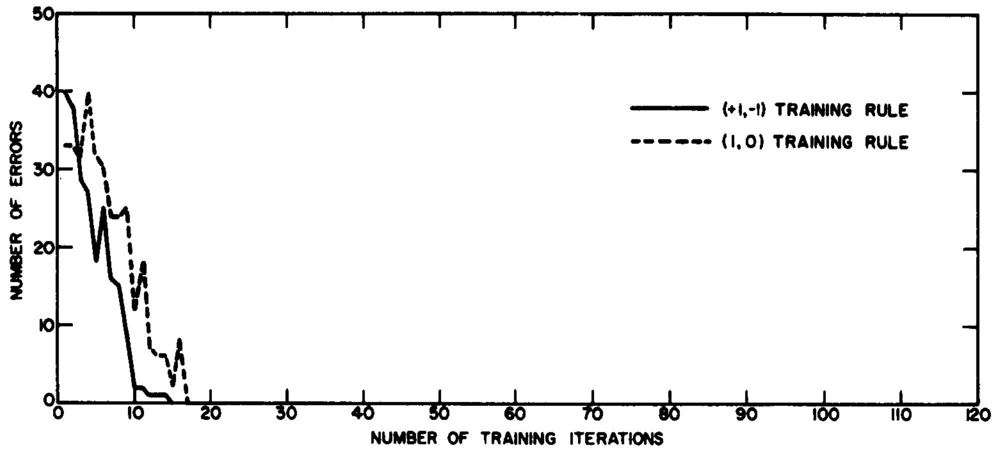

FIG. 2 LEARNING CURVES FOR PATTERN SET 1
(Each pattern containing exactly five ones)

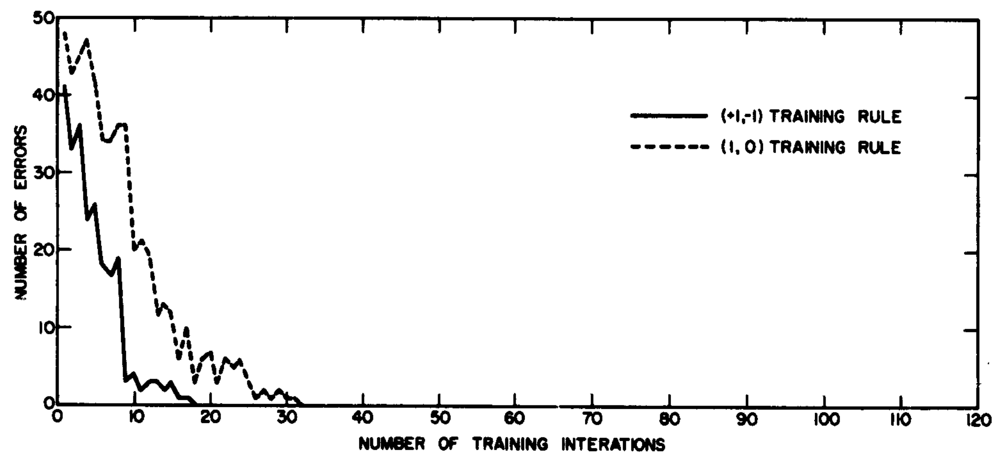

FIG. 3 LEARNING CURVES FOR PATTERN SET 2
(Each pattern containing exactly five ones)

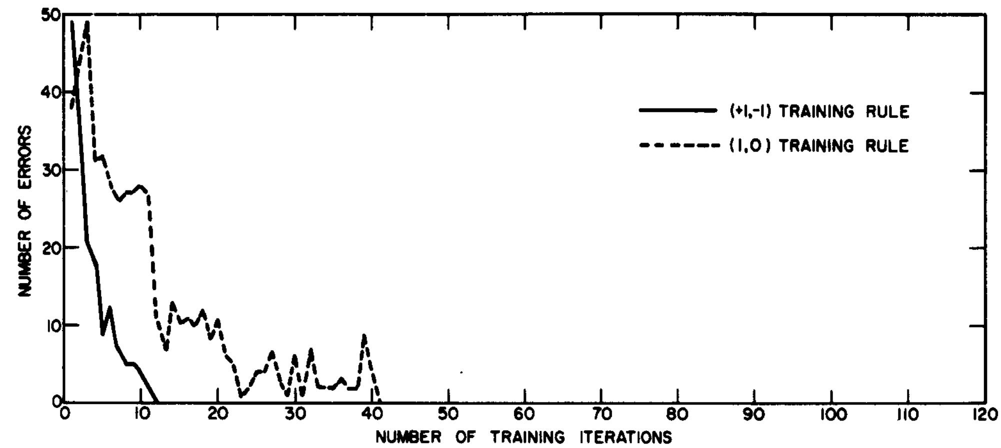

FIG. 4 LEARNING CURVES FOR PATTERN SET 3
(Each pattern containing exactly ten ones)

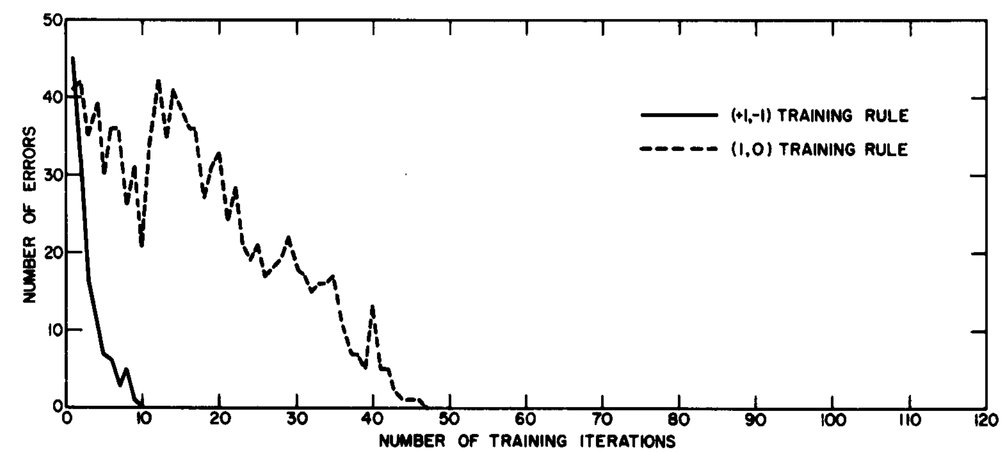

FIG. 5 LEARNING CURVES FOR PATTERN SET 4
(Each pattern containing exactly ten ones)

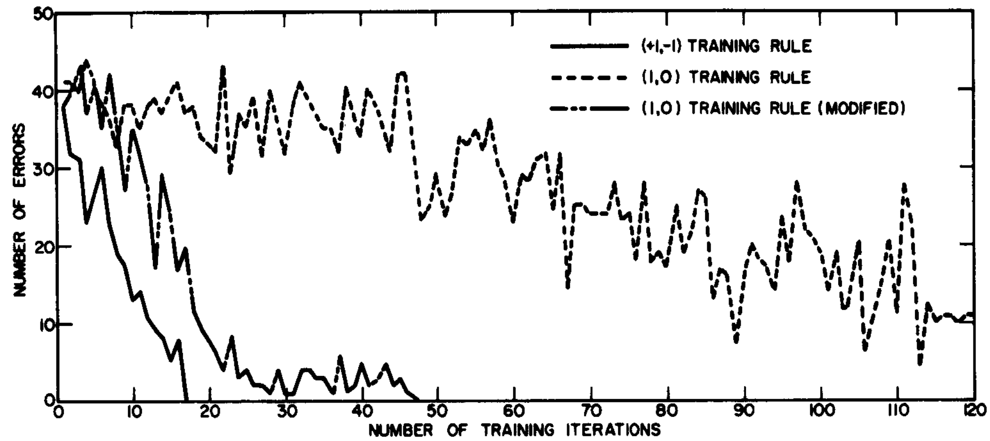

FIG. 6 LEARNING CURVES FOR PATTERN SET 5
(Each pattern containing exactly fifteen ones)

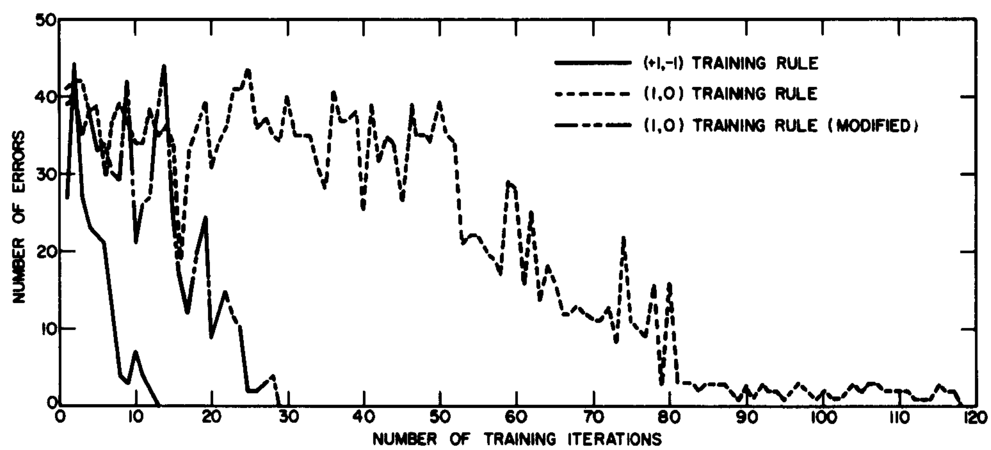

FIG. 7 LEARNING CURVES FOR PATTERN SET 6
(Each pattern containing exactly fifteen ones)

E. CONCLUSIONS

As a result of the above experiments, it has been concluded that the convergence of the (+ 1, - 1) rule is sufficiently faster than that of the (1, 0) rule to warrant the expense of the more complex circuitry needed to implement the (+ 1, - 1) rule. It has also been concluded that the (1, 0) rule can be modified, if desired, to effect a substantially faster convergence rate than the unmodified (1, 0) rule.

III TESTING TAPE-WOUND CORE PAIRS FOR UNIFORMITY

A. SUMMARY

The tested cores for MINOS II, which were supplied in four shipments by Magnetics, Inc., were sorted into four categories: "OK," "?," and "NG," representing progressively larger tolerances about the mean values of two parameters for each batch (shipment), and "WE," representing cores that were disqualified because of poor erasure. Approximately 57.5% were deemed "OK," 32.7% deemed "?," and 6.9% deemed "NG." The remainder of those tested (2.9%) were disqualified because of poor erasure. The cores were tested in pairs using a test program that simulated actual operating conditions. This test was in lieu of the more definitive (but more time-consuming) tests that were applied to several wired arrays of weights described in Sec. IV.

It was found that the variance among core pairs of core-pair parameters, as measured in the test circuit of Fig. 8, was too large for the twenty factory-tested pairs of cores to establish reliable mean values. It was also found that matching cores in pairs according to their major hysteresis loop does not substantially reduce the variance. A small sample of matched core pairs, selected from cores chosen at random from all four batches, had a percentage standard deviation of 21.4%, while a larger sample of unmatched core pairs from Batch 1 had a percentage standard deviation of 20.2% for one of the tested parameters.

B. INTRODUCTION

Because of the uncertainty involved in the construction of such a comparatively large and unique machine as MINOS II, it was thought prudent to test a majority of the tape-wound cores. A test was devised that simulated actual machine operation, yet remained simple to perform. The intent was to shorten the testing time per core pair so as to allow a large sample of core pairs to be tested. The results were to provide an indication of other operational parameters of each core pair, testing

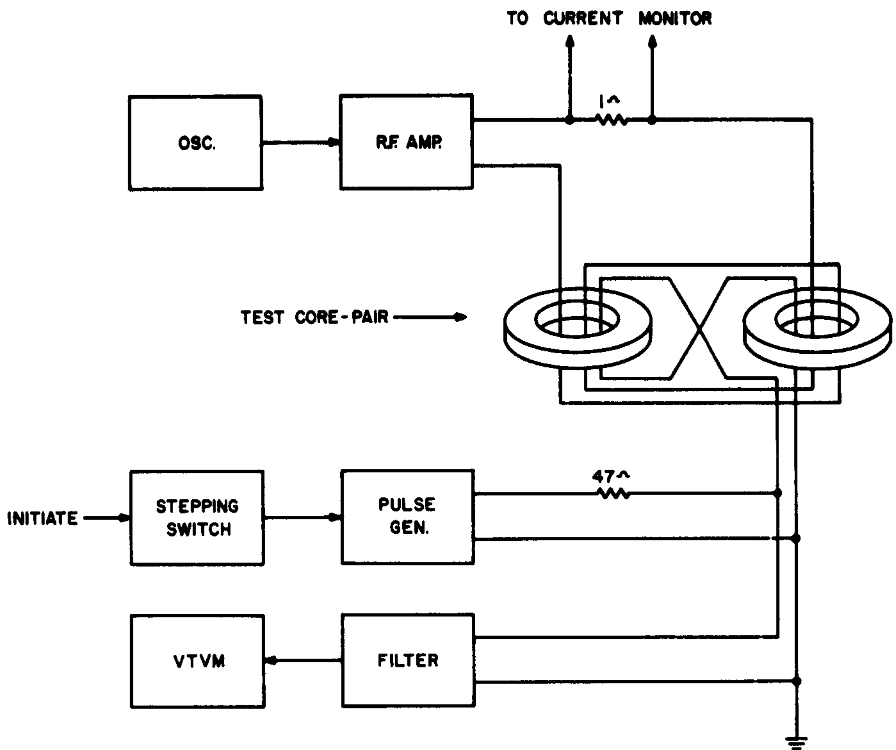

FIG. 8 CORE PAIR TEST CIRCUIT

of which would require more time-consuming measurements. These operational parameters are the minimum and maximum values, i.e., the end points of a range of adapt current at a fixed high-frequency drive current magnitude and frequency for which the stored value (remanent flux state) of the core pair will change in the presence of, and will not change in the absence of, the high-frequency drive current. The equivalence of the tests rely on the observed characteristic that the irreversible flux switching rate is roughly proportional to the magnetomotive force in excess of the threshold value. In the tests described in this section, a fixed value of adapt current was used to change the core states in combination with a fixed value of high-frequency drive current. For some core pairs, this value of adapt current would exceed threshold by a larger amount than for other core pairs, which would result in a larger change in the remanent flux state of the former than would be observed for the latter. A mean flux change was established by measuring the factory-tested cores as standards for each batch. The measurement was of the amplitude and phase of the second-harmonic pulses for a series of 50 short (50 μsec) applications of adapt current. Progressively increasing tolerance about the mean values determined the respective categories of "OK," "?," and "NG."

C. DETAILS OF THE TEST PROCEDURE

Each core pair was placed in a special test jig, which conveniently implements the circuit of Fig. 8. The core pair was first erased by 10-kc current through the high-frequency drive winding; the erase current was slowly reduced in magnitude to zero. The second-harmonic readout voltage in the erased state was usually less than three percent of its maximum possible value. (Any core pair that would not erase to less than 10% maximum value was placed in the "WE" category.) The erased state was chosen as a convenient reference point because the switching rate from saturation tended to depend on how hard the core pair was initially saturated.

A special stepping-switch circuit was built to supply 50 trigger pulses at each push of the INITIATE button. This was a time-saving convenience, whose extensive use more than repays the time required for its construction. Each trigger pulse initiated an adapt pulse whose amplitude, duration, and rise and fall times were controlled by dial settings on the pulse generator. A standard oscillator and ultra-low-distortion amplifier were used to supply the high-frequency drive current. Both the adapt pulse and the high-frequency drive current were constantly monitored.

The first measured parameter was the amount of second harmonic in the readout voltage after the application of 50 standard adapt pulses (200 ma-turns, 50 μsec) in the presence of high-frequency drive current (1.32 amp-turns peak-to-peak, 100 kc) with the core pair initially erased. The second measured parameter was the amount of second harmonic in the readout voltage after the application of a total of 250 adapt pulses (200 more, following the first 50), again with high-frequency drive current. The pulse amplitude and duration were chosen so that 50 pulses would switch the core pair to approximately one-third maximum value and 250 pulses would switch the core pair into saturation. Thus, the switching rate and maximum readout value were obtained for each core pair. Using the results of some preliminary testing for which the values of readout voltages were recorded, tolerance limits of ± 25% of the mean value of the 50-pulse readout voltages and ± 10% of the mean value of the 250-pulse readout voltages were placed in the "OK" category. Tolerance limits of ± 50% and ± 25%, respectively, for the 50- and 250-pulse reading of those core pairs not in the "OK" category defined the outer limits of the "?" category. All other core pairs were placed in either the "WE" or "NG" categories. If the two readings fell into different categories, then the wider tolerance category was chosen. If the category was "?," the measurement was repeated and a two-out-of-three choice was made.

D. RESULTS

Table I shows the results of testing 4209 core pairs in this manner, listed by batch number. Note that 67.2% of the core pairs tested from Batch 3 were classified "OK," while only 48.9% of the core pairs of Batch 4 tested "OK." While statistical variations are certain to exist from batch to batch, some of the variation in the percentage "OK" may be attributed to the variation of sampled mean values derived from an insufficient quantity of "standard" cores for each batch. The "standard" cores were chosen to be those that were factory-tested for total flux, erased flux, switching time, and coercive force. It was felt that any correlation between these latter parameters and those measured in our tests would be useful if, at some later date, it were necessary to either find causes of difficulty or to find a simpler test that could be easily implemented by the manufacturer without undue change in his equipment. Although a correlation analysis has not yet been made for these cores (because more important tasks are at hand) it is suggested that some effort be given this problem in the future. A rough inspection did not show any strong correlation between our measurements and those of the manufacturer either for the previously factory-tested cores or for a small batch of cores returned to the manufacturer for testing.

Table I

PERCENTAGE OF BATCH TOTAL OF "OK," "?," "NG," AND "WE" CORE PAIRS

Batch	Percentage "OK"	Percentage "?"	Percentage "NG"	Percentage "WE"	100% Equals
#1	60.0	33.8	6.2	0	943
#2	54.1	35.3	7.9	2.7	964
#3	67.2	28.2	4.6	0	1137
#4	48.9	33.6	8.7	8.8	1165
Averages of Total Tested	57.5%	32.7%	6.9%	2.9%	4209 core pairs

The manufacturer found a greater number of pairs matched according to major hysteresis loop among the "OK" and "?" than among the "NG" core pairs, and suggested that matching might increase the percentage of core pairs categorized as "OK." However, 48 matched pairs were compared to 547 unmatched pairs and the percentage standard deviation for the 50-pulse reading for the matched pairs was higher (21.4%) than for the unmatched pairs (20.2%). Two considerations detract from the strength of this result: First, 48 core pairs may be too small a sample to give a good measure of the actual variance, the matched core pairs were from all four batches, and the unmatched core pairs were chosen only from Batch 1. (See Figs. 9 and 10 for the matched and unmatched cases, respectively.) The assumption was made that the distribution of readings (percentage of total having a reading whose value is less than X) is Normal. As plotted on the coordinates used in Figs. 9 and 10, a truly Normal distribution would be a straight line. The straight line was plotted on the graph so as to produce the minimum apparent error. This graph is useful in that in a Normal distribution, approximately 16% on the abscissa corresponds to the mean value, $\overline{V}$, minus the standard deviation, σ, i.e., $\overline{V} - \sigma$, on the ordinate axis. Similarly, 50% on the abscissa corresponds to the mean value, $\overline{V}$, on the ordinate axis. Thus, σ and $\overline{V}$ are directly measured from the graphs.

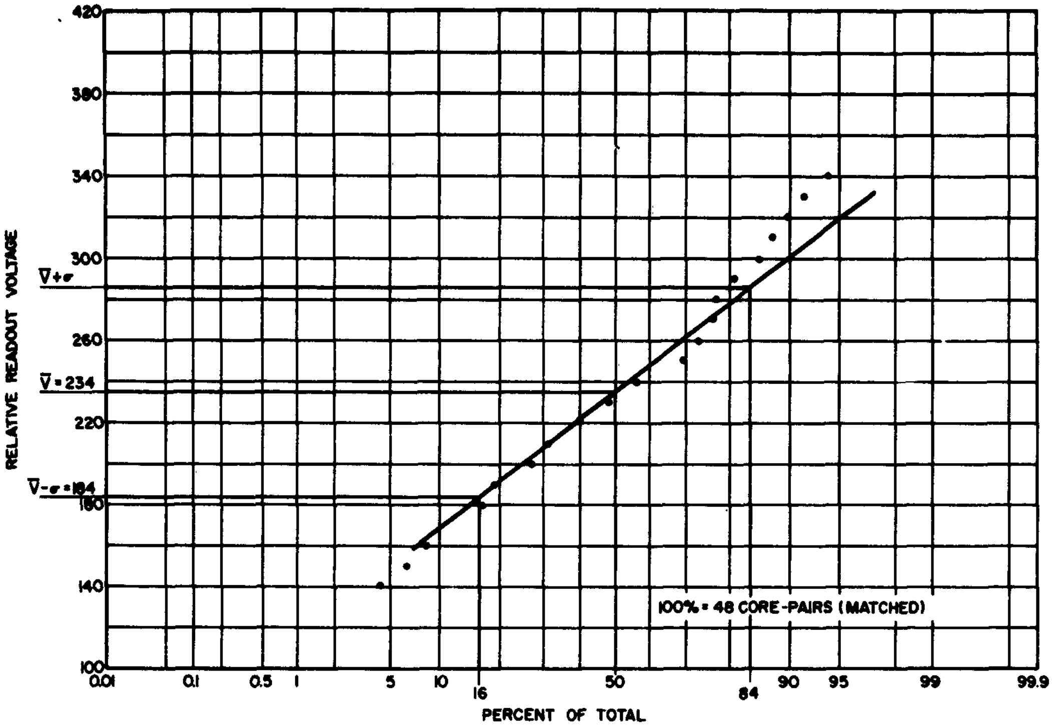

FIG. 9 DISTRIBUTION CURVE OF RELATIVE READOUT VOLTAGE OF MATCHED CORE PAIRS

The straight lines were chosen by eye so as to minimize the RMS error between the plotted points and the lines, by weighting those points representing small percentages from zero and 100 less heavily than those around 50%. The difference in the matched and unmatched means is most likely due to a change in a constant multiplication factor (e.g., meter calibration) and does not affect the calculation of the percentage standard deviations.

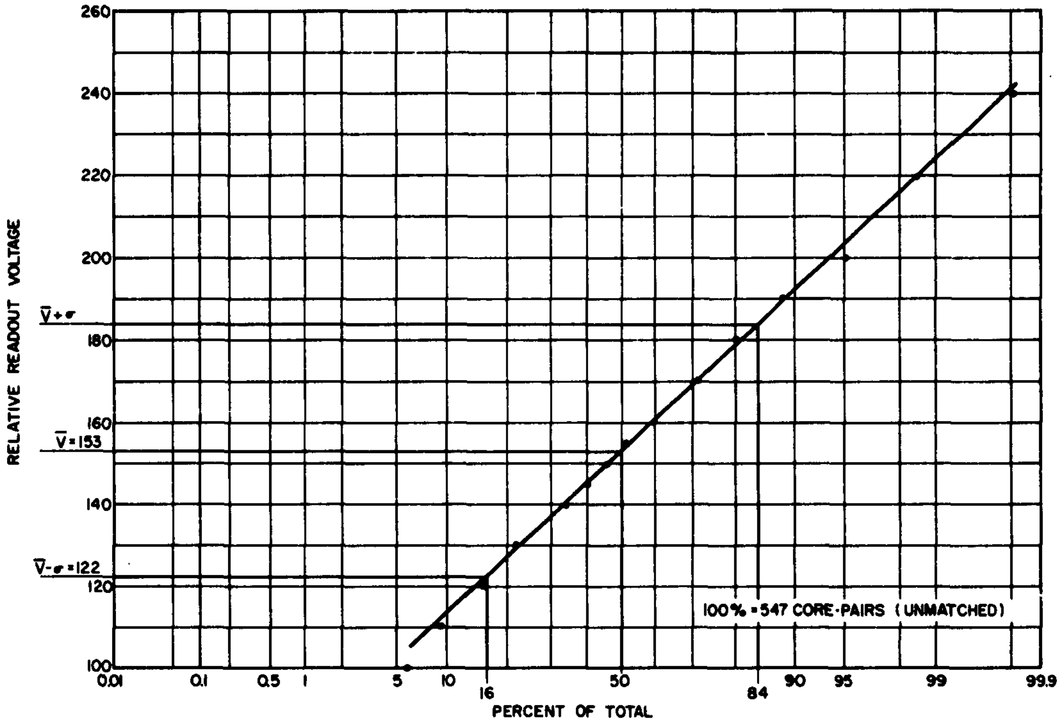

FIG. 10 DISTRIBUTION CURVE OF RELATIVE READOUT VOLTAGE OF UNMATCHED CORE PAIRS

IV CONSTRUCTING AND TESTING THE WEIGHT ARRAYS FOR MINOS II

A. SUMMARY

Eleven of the twelve weight arrays for MINOS II have been wired and three have been tested. All tested gave satisfactory results.

B. DESCRIPTION OF ARRAY CONSTRUCTION

The core pairs are supported by the current-carrying conductors. Each core is located in a plane at the intersection of an (input) high-frequency drive line and an (output) readout-adapt line, the latter being oriented perpendicularly to the former (see Fig. 11). Although single-turn circuitry would have been possible, each line in the present construction links the cores four times in order to economize in the drive and readout circuitry. Each high-frequency drive line threads, in the same sense, one core of each of the thirty-three core pairs linked by that line; the drive line returns through the remaining core of each core pair linked by that line in the opposite sense. Each readout-adapt line threads both cores of the seventeen core pairs linked by that line in the same sense and returns back externally to the cores. The return paths of the lines are as close to their respective forward paths as ease of construction would allow so to minimize electromagnetic cross-coupling (see Fig. 11 of QPR #10 for the schematic diagram).

Because connectors are usually among the least reliable components in electronic circuitry, our attempts to economize did not affect our choice of connectors. The use of connectors was deemed necessary to facilitate assembly and removal of the array. Since weight and space are not pertinent criteria for the construction of this experimental machine, each array is constructed as a rugged and well-protected unit. Aluminum extrusions were chosen for the frame, AWG 24 copper wire is used in stringing the arrays, and a layer of clear, rubber-like epoxy is placed over the cores and wire (not shown in Fig. 11).

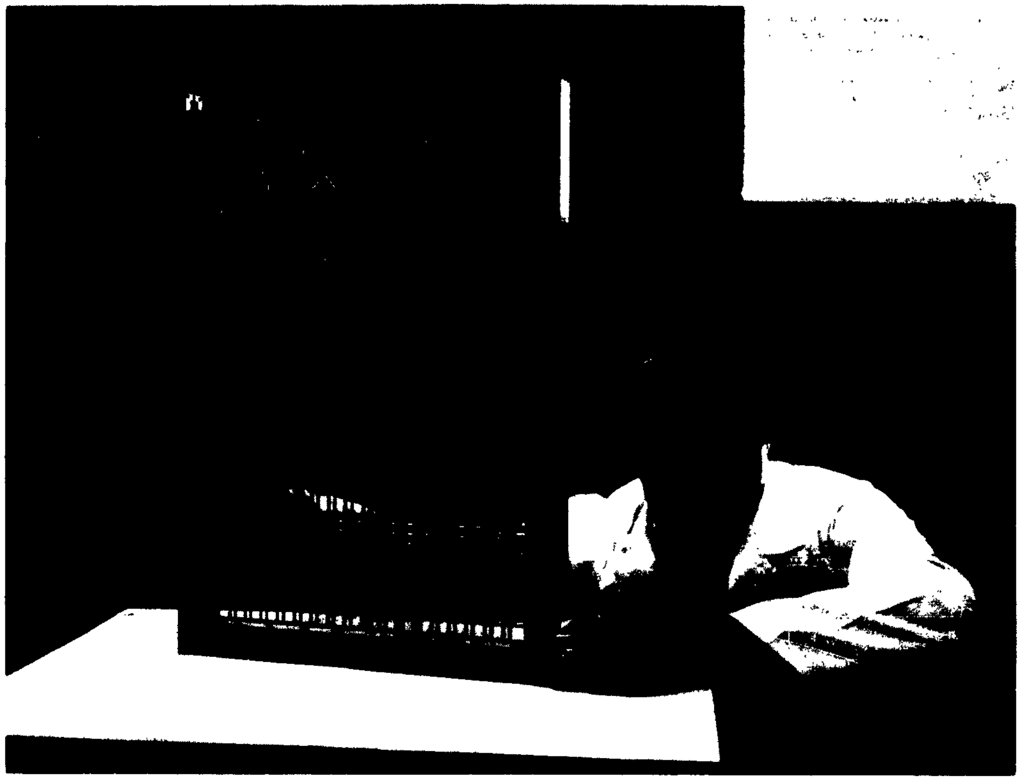

FIG. 11 A WIRED CORE PAIR ARRAY

Provision has been made for a circuit board on which the transistor gate circuitry will be mounted. The gate circuitry connects any desired combination of high-frequency drive lines to a regulated 100 kc voltage source and is controlled by dc gating signals.

C. WEIGHT ARRAY TESTS

The weight arrays were tested to find the usable range of adapt current for each readout line. The usable range of adapt current is defined as that region of current in which all weights (on the readout-adapt line being tested) will change their stored value (change the net remanent flux) in the presence of a high-frequency drive current and no weight will change its stored value in the absence of a high-frequency drive current. The results of the tests to date are shown in Fig. 12.

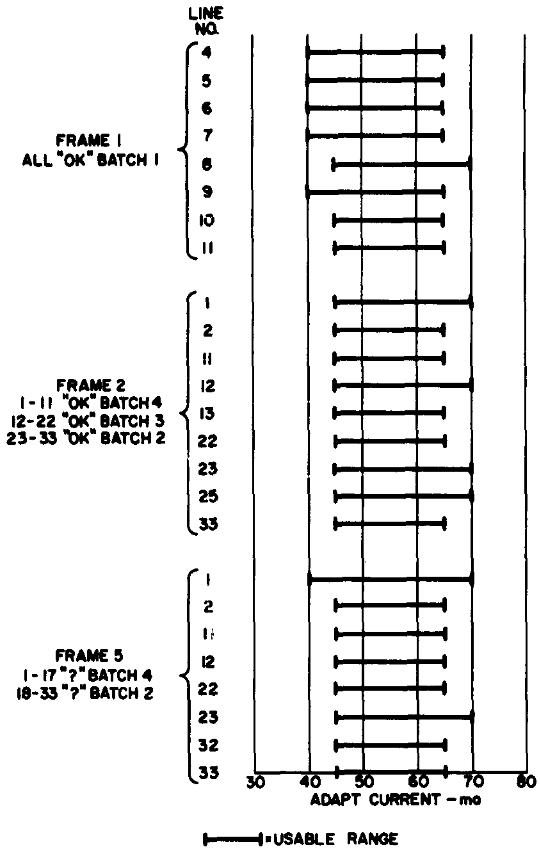

FIG. 12 USABLE RANGES OF ADAPT CURRENT FOR TESTED READOUT-ADAPT LINES

As mentioned in Sec. III, the cores were shipped in four batches, and each batch was tested according to the mean values of the factory-tested parameters of the cores from that batch. Accordingly, the tested core-pairs were labeled, "OK," "?," or "NG," corresponding to respectively increasing tolerance limits about the mean values for each batch. Array #1 used all "OK" cores from Batch 1. Array #2 used all "OK" cores from

Batches 2, 3, and 4. Array #5 (the third to be tested) used "?" cores from Batches 2 and 4.

All core-pairs were tested at a fixed value (330 ma. peak-to-peak) of high-frequency drive current, which had been previously determined (Sec. III, QPR #10) to be optimum for these cores. With a high-frequency drive current applied, fifty pulses would change the stored value by more than one percent of the maximum stored value, while a total of more than one thousand pulses would not change the stored value by more than one percent of the maximum in the absence of high-frequency drive current.

The test results show a usable range from 45 ma to 65 ma (four turns) on all three arrays, even though the tolerances on the core-pairs were different according to arrays. On the average, the arrays with "OK" cores had a slightly larger range, but were shifted higher or lower according to batch (see Fig. 12).

The above tests used approximately six percent of all weights. The tested readout-adapt lines were chosen from differing categories of core-pairs. Although we feel that all of the weight arrays should operate as planned, we are continuing to check the arrays as soon after assembly as time permits.

V SUPERVISORY TRAINING SYSTEM (TEACHER) FOR MINOS II

A. GENERAL ORGANIZATION OF THE MACHINE

In order to discuss the supervisory scheme in detail, it is necessary to describe the over-all machine organization, the functioning of the logical elements, and the training rule to be used in adapting the weights to perform useful tasks on pattern recognition and related problems. A simplified block diagram of the machine is shown in Fig. 13.

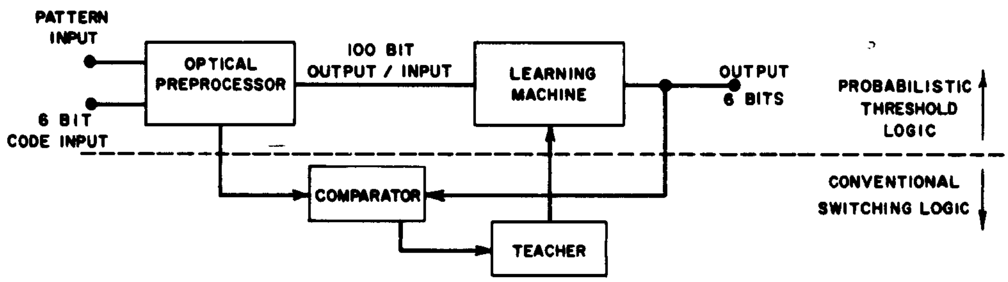

FIG. 13 BLOCK DIAGRAM OF MINOS II

Some aspects of the preprocessor design were discussed in Quarterly Progress Report 10. This part of the machine is not adaptive on a short-term basis, although a new set of masks can be inserted in the preprocessor effectively to rewire the input to give optimum performance on a specified limited class of patterns. Figure 13 shows that the preprocessor output feeds the learning machine input and that the comparator provides the supervisory and control logic, which performs the teaching operation; signals controlling the incrementing and decrementing of weights are derived by comparing the actual machine output and the desired output code; the latter may be specified optically in the preprocessor, together with the pattern.

Figure 14 is a schematic representation of the machine and shows the interconnections between the various sections of the threshold logic elements without considering any of the supervisory or control functions. It has been decided, for reasons which are demonstrated in Sec. II of this report, to construct a machine capable of implementing the (+ 1, - 1) training rule, rather than only the (+ 1, 0) training rule, which has serious limitations.

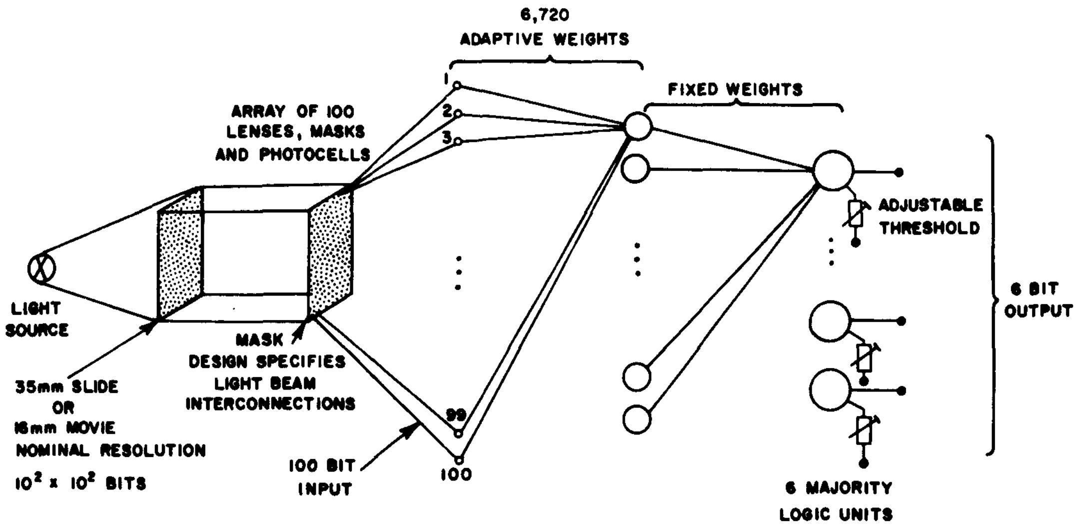

FIG. 14 THRESHOLD LOGIC INTERCONNECTIONS

B. COMPARISON OF INPUT AND OUTPUT CODES

Because it is intended to make the new machine as automatic as possible, the training logic must be built in as an integral part of the structure. Each training input pattern, whether it be a high-resolution graphical image fed in via the slide or movie projector, or a manual input via the touch-sensitive retina, will be associated with a six- or nine-bit classification code, which specifies the desired output of the machine.

The comparator compares the input and output codes, and two broad possibilities arise. If the machine output is correctly classified, the comparator output must give the appropriate signal and arrange that no alteration of the weight values is made. If the output code is in error, the incorrect majority logic units must be identified and these must enter the adapt phase of the training cycle. For the 6-bit output code arrangement, each majority logic unit is fed by 11 inputs, and the majority logic training rule requires (a) that the minimum number of inputs that will make the majority correct must be trained, and (b) that the ones trained must be those closest to threshold. This training rule (and digital computer simulations based upon it) has been described in more detail in earlier reports. The sequence of principal operations required to carry out the training logic is as follows:

(1) Pattern and classification code are presented simultaneously. If the pattern is manually set up, the INITIATE push-button must be touched to start the training cycle. The INITIATE button is the upper left hand coding element in the operator's retina--cell number 101. If the pattern is presented by means of a slide, the "initiate" coding photocell is allowed to receive light, and thus the silicon controlled rectifier (SCR) connected to it will fire, initiating the cycle. This defines time zero.

    When presenting a pattern manually on the touch-operated retina, the dc supply to the SCRs must be on in order to light the lamps so that the operator can see the pattern he has written. This means that the inputs to the learning machine will vary as the pattern is written up. Also, the final outputs are then free to vary and operate the output display. This will probably be a useful feature, allowing the operator to compare input and output patterns while altering the inputs. The initiate coding cell then assumes the function of triggering the automatic training cycle, with adaptations if necessary. The simulated adaptation using the ramp signal will therefore not commence until just after the initiate cell is on.

(2) Comparator presents a binary output code word, which may have some incorrect bits.

(3) For the incorrect bits, the majority logic units that are to be trained enter the adapt phase.

(4) The simulated adaptation is carried out to determine exactly how many and which of the eleven inputs to the incorrect majority logic units need be changed.

(5) The adapt phase is carried out. All necessary increments and decrements are made to the weights.

(6) The adapt phase is completed, and the next pattern presented, thus completing one cycle.

It must be decided whether the adaptation is to be made on a fixed-time basis or on the basis of a correct output, i.e., the end of the adapt phase must be determined either by a fixed time interval after initiation or by a correct output. Previous research has shown that convergence is more rapid when incorrect decisions are always corrected before proceeding to the next pattern than if a reasonable fixed increment is made for all incorrect decisions, regardless of whether this fixed increment is the appropriate size to correct the pattern classification. Either process will converge and the difference in efficiency appears to be small. However, to eliminate out-of-step control signals, we must ensure that the adapt pulses occur only within some prescribed interval. An adaptation of appropriate size may be made in less than 2 msec; in general, errors will be of such a size that about five increment pulses will correct most patterns. If the slide is left on for 0.9 sec., a large number of adaptations will be made for each pattern, and most patterns are sure to be corrected. However, the next pattern will be presented only after the fixed adapt cycle. Note that this is compatible with a movie projector system in which the sequence of presentation of patterns and their timing is not under external control, i.e., no patterns can be presented out of sequence and the instants of presentation of the patterns cannot be synchronized to any fixed clock cycle because of the speed variations in the projector motor. This synchronization is also not possible with the 35 mm. slide projector. It is, therefore, advisable to control the machine logic using timing pulses from the projector. However, it must be arranged that the projector may only make changes when these are acceptable from the machine's point of view, i.e., after adaptation has ceased, and this can be ensured with both projectors by continuing the adaptation for a limited fixed time, shorter than the presentation period. For the purposes of discussing the control scheme, we will consider operation with the 35 mm. slide projector only. The continuous run-through speed for this projector (Kodak Carousel) is approximately 1 slide per second, and the duty cycle is approximately 0.25 second on and 0.75 second off. (The movie projector duty cycle is shown in Fig. 15 and is included here for future reference.)

The time taken for an adaptation, plus the associated resettling of the readout amplifier, flip-flops, etc. to new values is not greater than 2 msec. The simulated adaptation requires a "ramp" signal to be added; this will take approximately 10 msec. After the initial preprocessor adjustment, the first simulated adaptation may begin and the adaptations may take place in succession thereafter until either the output becomes correct, or the next pattern is due to be presented, i.e., after about 0.8 sec. It appears logical to allow the slide projector to run continuously at its own speed and not stop until convergence is obtained. This is compatible with movie operation, although in this case the time available for adaptation would be much shorter. The 35 mm. projector has a shutter that moves horizontally across the plane of the slide; and it is necessary to ensure that the coding element is on the side that receives the light last. A further 10 msec delay must be allowed to ensure reasonable mechanical stability. Thus, the output of the initiate SCR will first trigger a 10 msec one-shot delay circuit. A

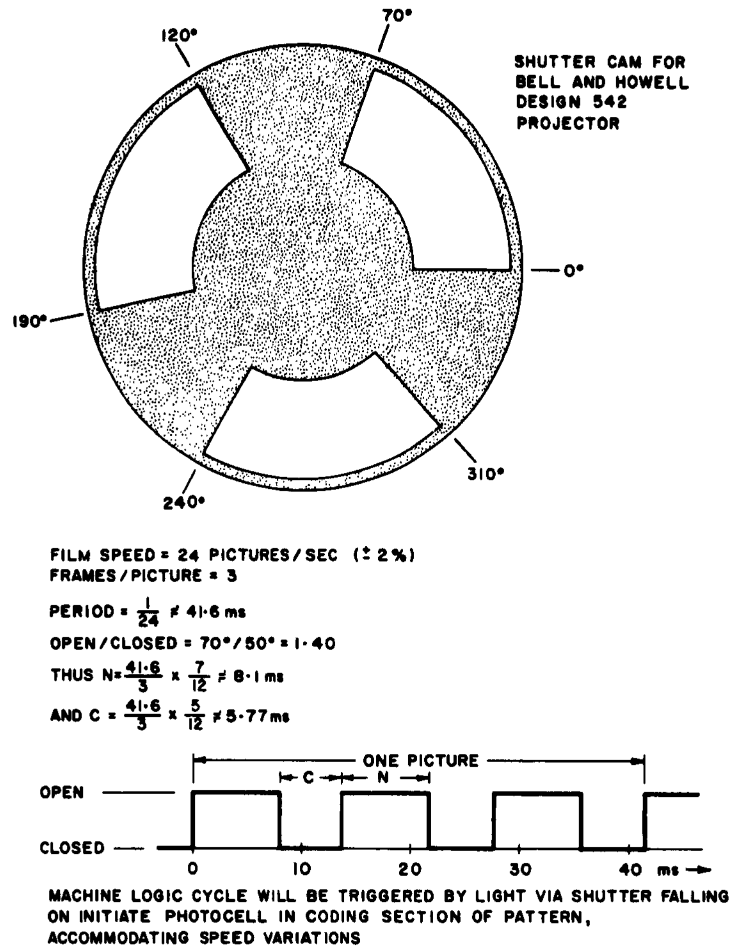

FIG. 15 PROJECTOR SHUTTER TIMING

delay of about 30 msec will now be allowed for the automatic level control to reach the value at which approximately half the preprocessor outputs come on. This fixed delay can also be obtained by a one-shot circuit. At the end of this time, say 40 msec, the preprocessor outputs will be stable and after a further 1 msec (approximately), the read-out amplifiers connected to the outputs of the weights will also have stabilized. Thus, after 41 msec, a reading from the learning machine output will be available. The comparator logic circuitry may have a potential repetition rate of 100 kc, so that the delay through the comparator is negligible. Thus, simulated adaptation can safely take place 42 msec after presentation of the slide.

C. THE 6-BIT COMPARATOR

Let the desired input code be represented by $f_i = (x_1, x_2, \ldots, x_6)$ and the output code by $f_o = (y_1, y_2, \ldots, y_6)$. Each x and y represents a binary valued variable, being the input/output code for the machine. An output, indicating the need for training, will only be given for cases in which the corresponding x and y terms disagree. The truth table for each bit is as follows:

Table II

"EXCLUSIVE OR" FUNCTION

Input x_j	Output y_j	f_j	Action
-1	-1	+1	No adaptation
-1	+1	-1	Decrement
+1	-1	-1	Increment
+1	+1	+1	No adaptation

The output of the jth comparator bit may be represented by the Boolean equation, $f_j = \bar{x}_j y_j \vee x_j \bar{y}_j$ which is the "Exclusive OR" function. The comparator identifies the majority logic outputs that are incorrect, and since these six are mutually independent and identical, the adapt logic circuit is repeated for each bit of the output code. Some possible methods to implement this "Exclusive OR" function are shown in Fig. 16. The actual circuit modules and true and false logic voltages have yet to be decided upon, bearing in mind such factors as compatibility with threshold logic voltage levels, price and availability of modules to be purchased, and production costs and availability of labor for those modules and circuits still to be manufactured. Each majority logic unit has 11 inputs from the outputs of 11 TLUs, which are in turn connected to 101 x 11

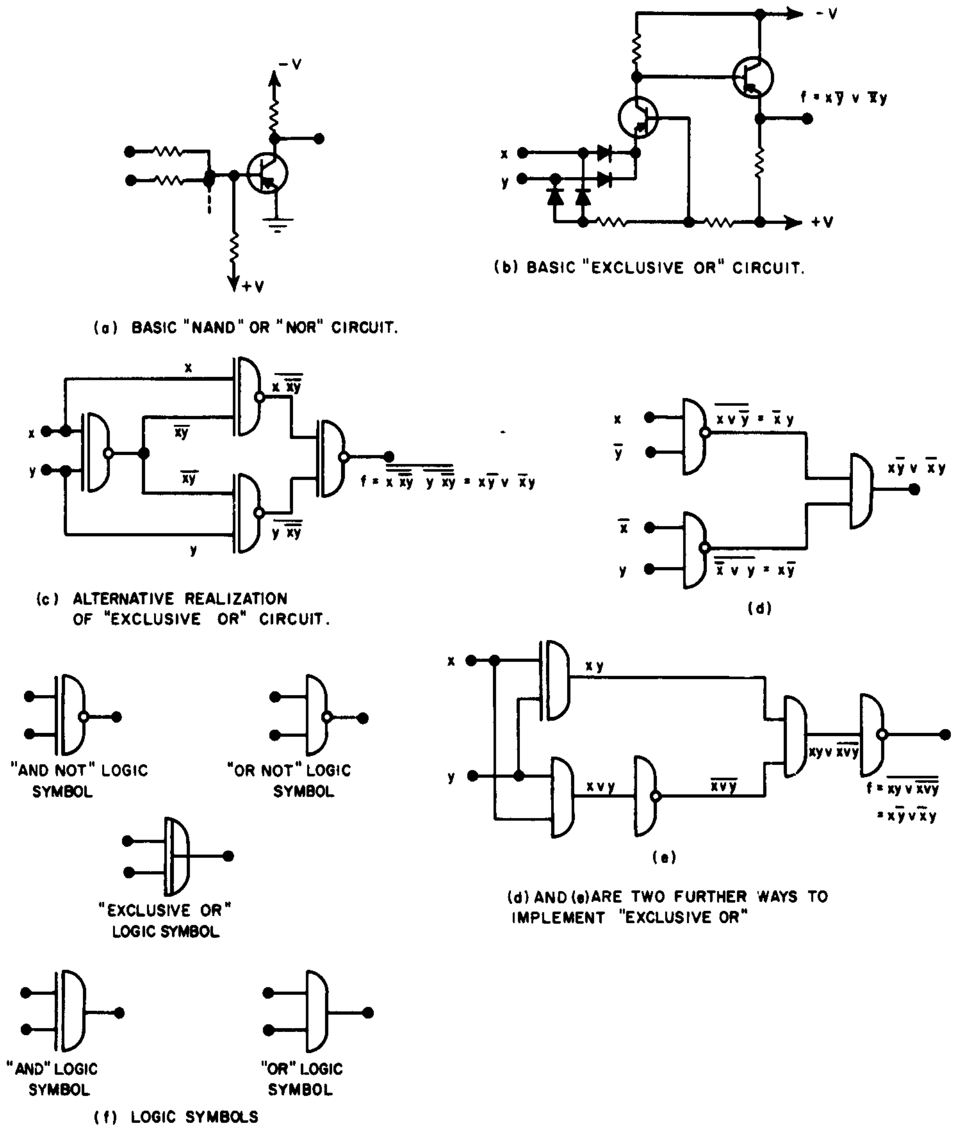

FIG. 16 "EXCLUSIVE OR" CIRCUIT FOR COMPARATOR
(a) Basic AND-NOT or OR-NOT Circuitry
(b) Realization of Exclusive OR Using AND-NOT Circuitry
(c), (d), and (e) Other Realizations of Exclusive OR
(f) Logic Notation

weights to be adapted. The design of this adapt logic circuitry will be considered in the next section, but may be regarded as an independent subsystem working within the supervisory system. The completion of the full cycle of the input-pattern logic sequence will now be considered.

D. ADAPTATION LOGIC SCHEME

Having identified the incorrect output bits as described above, the majority logic training rule must be implemented to adapt the weights. Each majority logic unit is of the form shown in Fig. 17, having 11 equally weighted inputs and a threshold of, say, 5.5. This logic element implements the Boolean function f_m, indicated algebraically. If the output is incorrect, it must be reversed, which defines the direction of training. The majority rule scheme we wish to implement requires identification of a certain number of threshold logic units (TLUs) with the lowest analog levels, for training. The TLUs having the lowest outputs may be discovered by simulating the adaptation and noting those units which change their response. A number of possibilities arise, as shown below:

Table III

POSSIBLE CONDITIONS FOR INCORRECT OUTPUT UNITS

NUMBER OF TLUs		
Wrong	Right	To Train
11	0	6
10	1	5
9	2	4
8	3	3
7	4	2
6	5	1

The simulated training may take the form of a signal added in series with the summed analog output of the weights, as shown in Fig. 18. The signal "ramp" must start at zero and rise in magnitude, its sign being determined by the desired direction of change of the output. The ramp signal must be added before amplification, since the amplifier is not

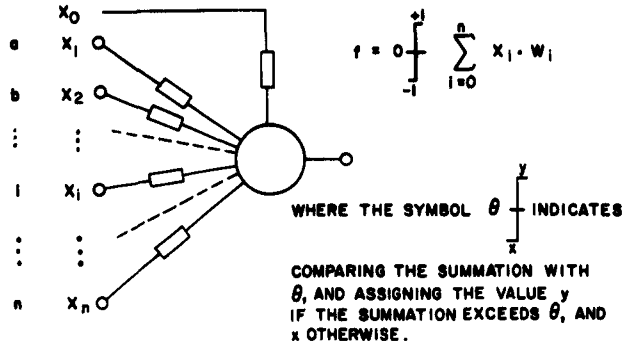

(a) GENERAL FORM OF THRESHOLD LOGIC UNIT.

FOR THE MAJORITY LOGIC UNIT, ALL $W_i = 1$ $(i \neq o)$
AND $W_O \cong \frac{n}{2}$, $X_O = -1$. THE BOOLEAN FUNCTION
IMPLEMENTED IS, FOR $n = 11$,
fm = abcdef v abcdeg v abcdeh v ···
THE ELEMENT TO BE USED IN MINOS II WILL HAVE
A VARIABLE VALUE OF W_O, IMPLEMENTING A MORE
GENERAL FORM OF MAJORITY, i.e., QUORUM.

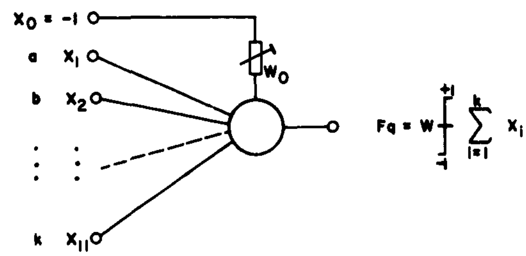

(b) QUORUM LOGIC UNIT FOR MINOS II

FIG. 17 MAJORITY LOGIC ELEMENT

linear and saturates for high values of input signal. This saturation or limiting of the amplifier does not affect the correct identification of weighted threshold logic units to be trained. As the ramp rises in magnitude, the TLU closest to threshold, but on the wrong side, will change the sign of its output, and this will also change the analog value of signal level controlling the classification. The quantized output of the majority logic element will only change if its analog value changes through the threshold, and this output bit will become correct at this instant. However, the output bit will only be correct if the majority of inputs to it is correct, and this majority will be exactly six, since the output was previously wrong, and the simulated adaptation was made

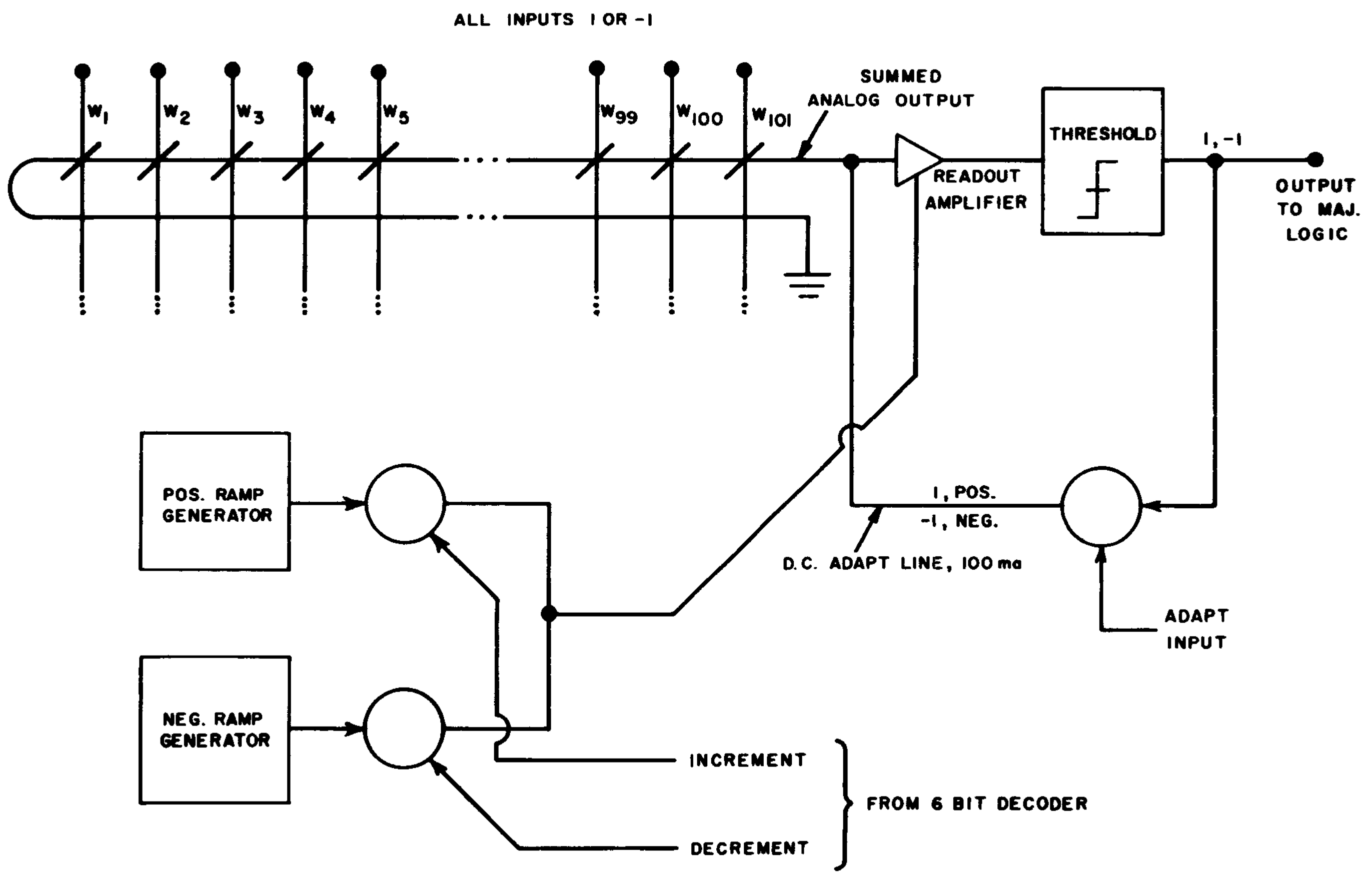

FIG. 18 ONE OF 66 ADAPTIVE TLUs

gradually. This procedure will be successful with any number of wrong voters (TLUs). The TLUs whose outputs changed during the simulated adaptation must be remembered, say by means of a flip-flop; the flip-flop may be used to gate-on the adapt pulse after the ramp has completed its full sweep. It is necessary to take the ramp up to a high value, since the wrong voters may all be saturated in the wrong direction. In some cases, because the analog values of several TLUs may be very close together, and because it takes a finite time to switch off the ramp, more than the required number of TLUs may be trained. However, this is not typical, and may be made less likely by decreasing the rate of rise of the ramp relative to the speed of the ramp switch-off circuits. The ramps to each of the six majority logic units will need to be switched off individually as their outputs become correct. Note that the comparator output will also automatically change; this indicates that the comparator must control the ramp circuitry.

The main part of the adaptive section of MINOS II is shown functionally in Fig. 19, for one of the 66 weighted input threshold logic units and for one of the output majority logic units. It should be realized, when reading this diagram, that certain sections shown have to be replicated 66 times, other sections six or nine times depending upon the number of ··t bits, and there is an over-all control section, only one of which .s required. In addition, the input gating circuits have to be replicated 200 times, since the memory system is divided into two parts, and there are 100 inputs to each section.

To summarize the adaptation logic scheme: The six-bit code comparator provides an increment, decrement, or no-train output for each of the six bits. This refers to the sign of the required change (if any) in the jth-bit, and determines the sign of the ramp signal to be added to all eleven summers contributing to the jth-bit inputs. Thus, the sign of the ramp signal, i.e., either in phase or out of phase, to be added to all eleven TLUs is determined. Now from these eleven, the smallest number of wrong voters must be selected to make the majority just correct. At the end of the simulated adaptation, this unique selection

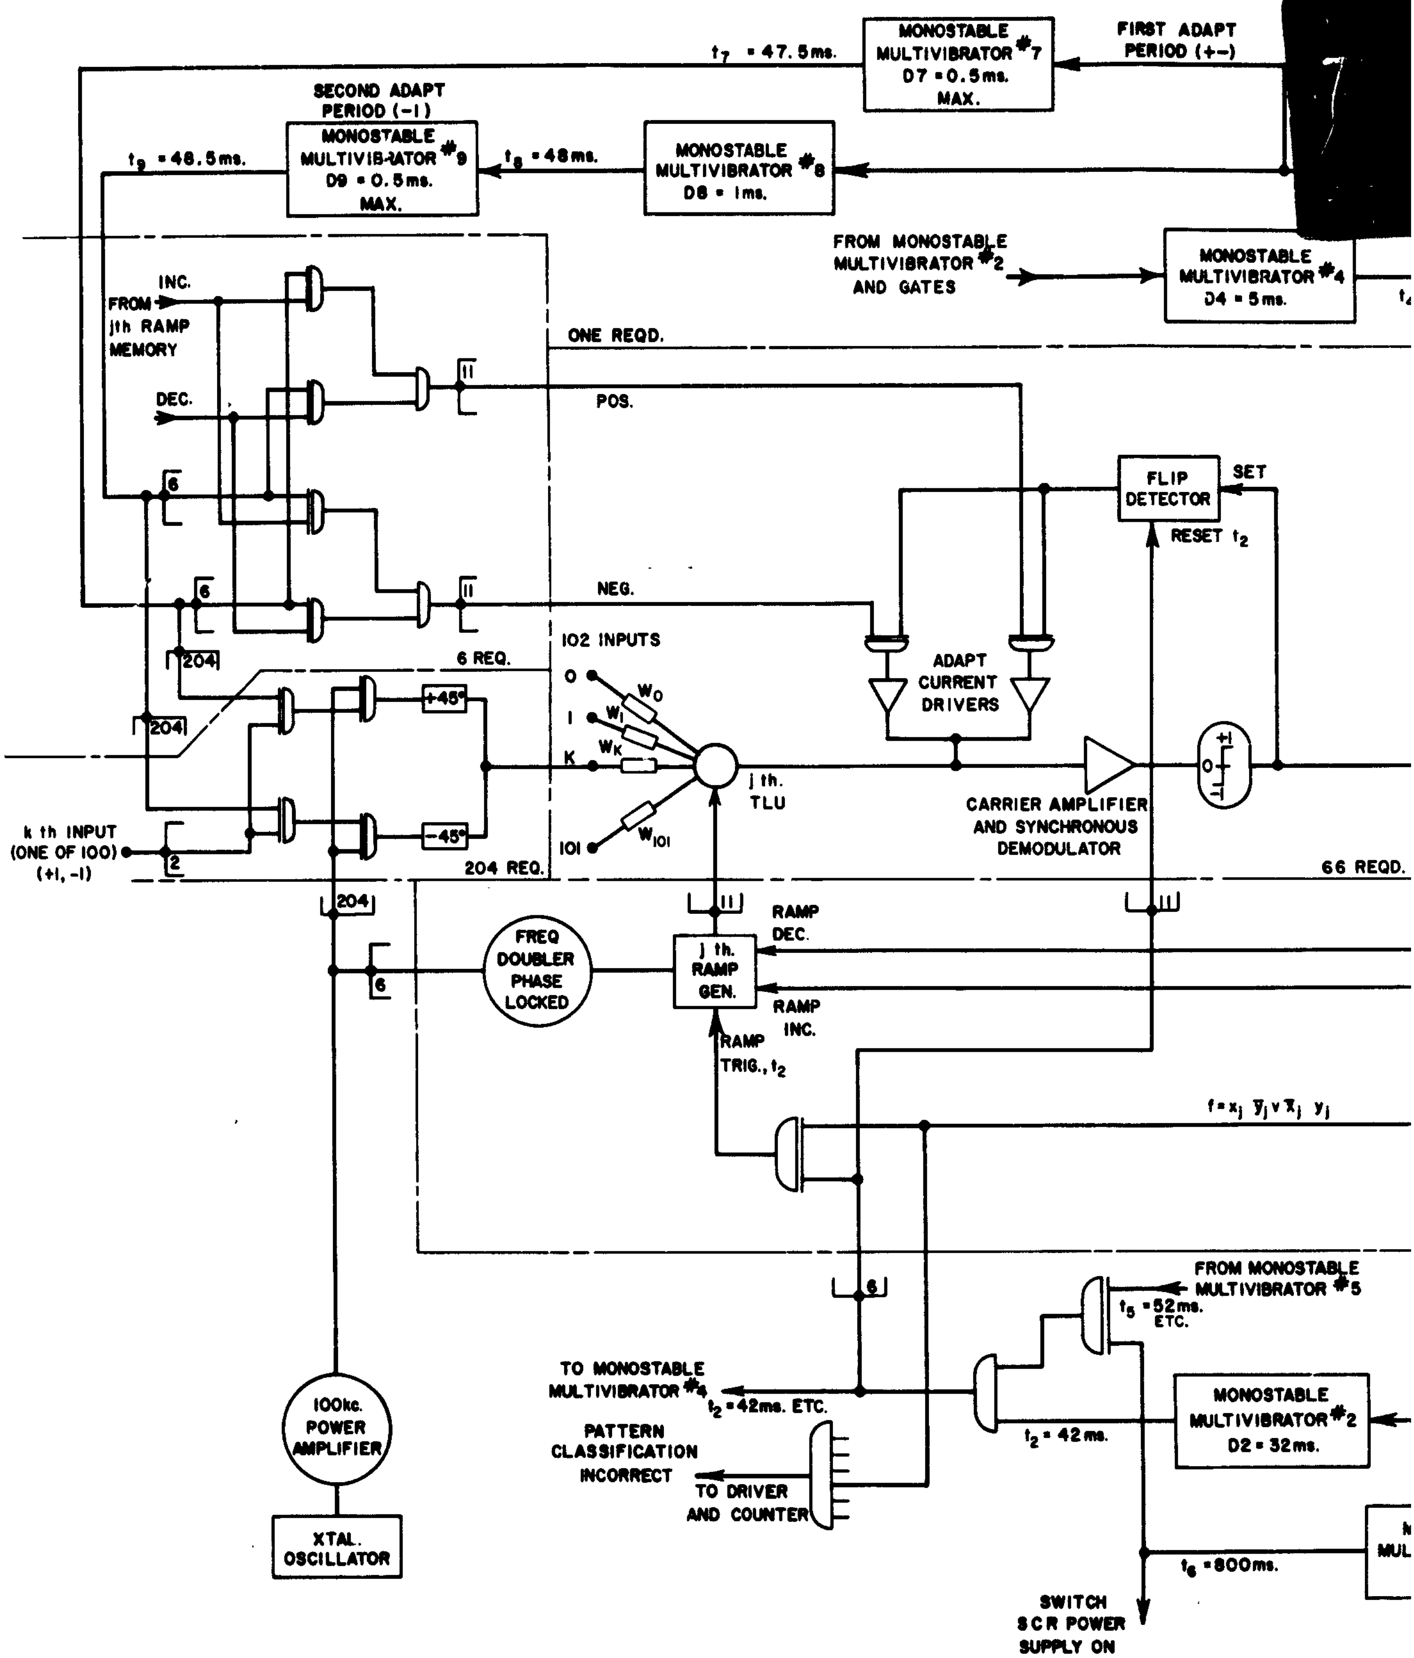

FIG. 19 AUTOMATIC TRAINING SYSTEM

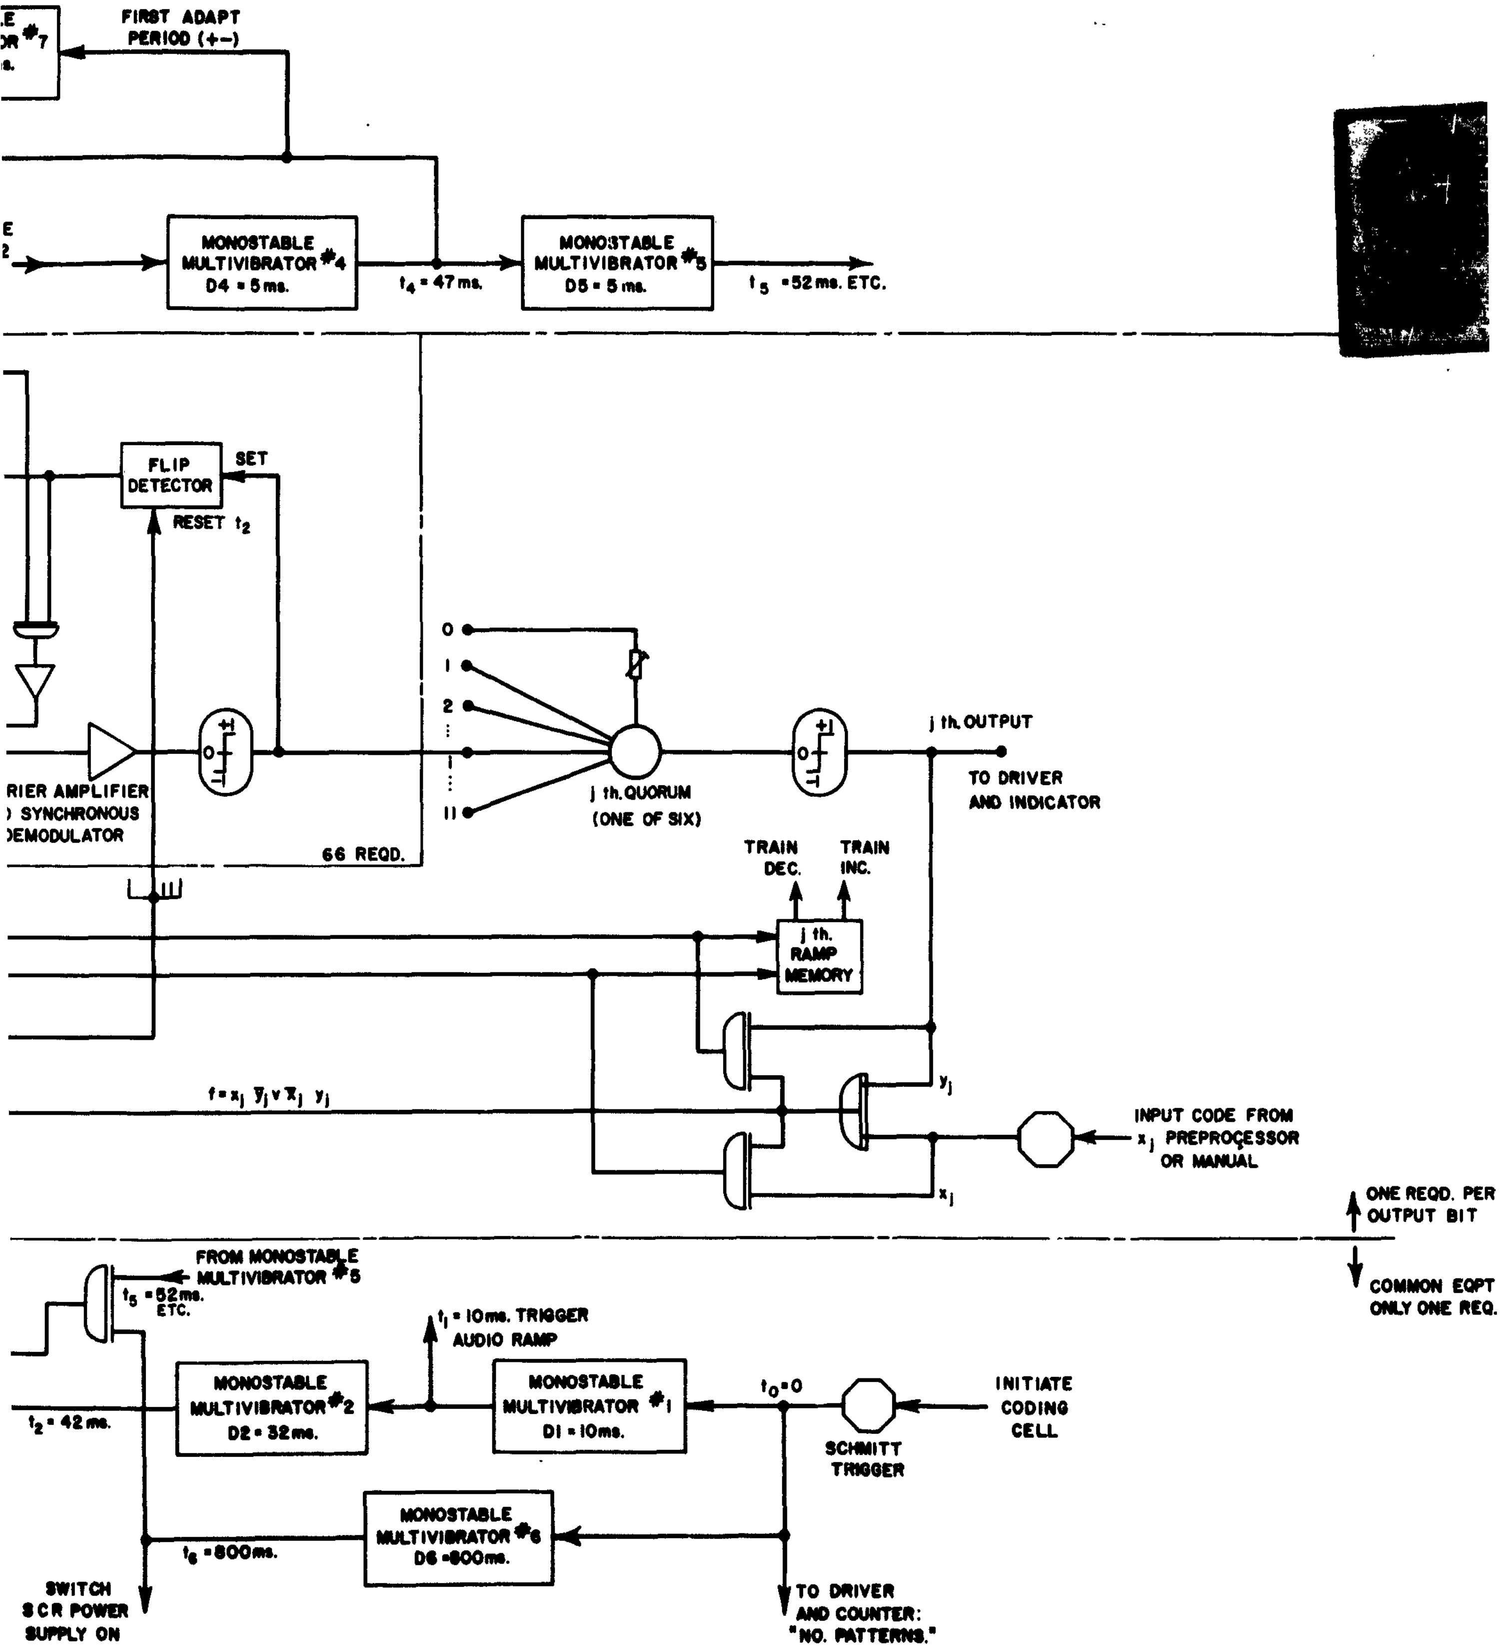

19 AUTOMATIC TRAINING SYSTEM FOR MINOS II

TRUTH TABLE FOR i^{th} INPUT TO j^{th} QUORUM

QUORUM OUTPUT y_j	DESIRED OUTPUT x_j	MAJORITY TRAIN RULE	INCREASE QUORUM SUM	DECREASE QUORUM SUM
−1	−1	−1	−1	−1
−1	−1	+1	−1	−1
−1	+1	−1	−1	−1
−1	+1	+1	+1	−1
+1	−1	−1	−1	−1
+1	−1	+1	−1	+1
+1	+1	−1	−1	−1
+1	+1	+1	−1	−1

TRUTH TABLE FOR TRAINING k^{th} WEIGHT ON i^{th} TLU

TLU TRAIN DIRECTION	INPUT SIGN k^{th} wt.	DC ADAPT CURRENT
DEC −1	−1	POS
DEC −1	+1	NEG
INC +1	−1	NEG
INC +1	+1	POS
DEC OR INC	0	0
NO TRAIN	+1 OR −1	0

FIG. 19 SHEET 2 OF 2

will have been made and those selected for training identify sets of 101 weights to be adapted. The sign of the code comparator, which defines the required change in the _j_th output bit (majority logic unit), also defines the required sign change in the output of the _i_th-weighted threshold logic unit. The weight connected to the _k_th input on the _i_th threshold logic unit must be adapted to change the output in the required direction defined by the output of the threshold logic unit, the code comparator, and the sign of the _k_th input. The _i_th-weighted threshold logic unit is typical of all 66, and is shown in Fig. 18. The threshold of the majority logic element is adjustable, so that values other than the five-to-six majority can be manually set. We have called this _quorum logic_ to denote a more general form of majority logic and to indicate that the threshold is variable. By varying the threshold of the quorum logic unit, its logical function can range from an "OR gate" to an "AND gate," with the majority function half way in between. The quorum logic unit is more particular than the threshold logic unit, since its input weights are all identical.

E. INPUT GATING AND CONTROL LOGIC

The truth table for training the _k_th weight on the _i_th-weighted quorum input, w_k is given in Fig. 19, and the implementation of this truth table will now be discussed. Changing the value of a weight in the memory plane is accomplished by the coincident presence of a direct current in the output line threading the weight concerned, and carrier current in the input line threading the weight. The direction of the change, i.e., increment or decrement, is determined only by the sign of the direct current and is unaffected by the phase of the carrier. This characteristic of the weight system as it is used at present requires that both positive and negative adapt current be applied to the output line to implement the truth table referred to above. This implies that two separate adapt phases are needed, one for positive adapt current and the other for negative adapt current. The weights may be similarly divided into two classes, those which require positive adapt current and those which require negative adapt current, according to the sign of their inputs and the required training direction. The inputs are either + 1 or - 1 and if these are turned on at separate times as shown in Fig. 19, the required logic scheme is implemented.

PROGRAM FOR THE PERIOD 1 MARCH TO 31 MAY 1963

All available time and effort will be devoted to completing the construction and testing of MINOS II.

IDENTIFICATION OF KEY TECHNICAL PERSONNEL
(For the Period 1 December 1962 to 28 February 1963)

	Hours Charged to the Project
Dr. Alfred E. Brain	
Research Physicist, Applied Physics Laboratory (Project Leader)	421
Mr. Harold S. Crafts	
Research Engineer, Applied Physics Laboratory	18
Mr. George E. Forsen	
Research Engineer, Applied Physics Laboratory	262
Mr. David J. Hall	
Research Engineer, Applied Physics Laboratory	291
Mr. Jack W. Machanik	
Research Engineer, Systems Engineering Department	8
Dr. Charles A. Rosen	
Manager, Applied Physics Laboratory	15
Shop and Technician	1,620
Clerical and Secretarial	41

DISTRIBUTION LIST

Organization	No. of Copies
OASD (R&E) Room 3E1065 The Pentagon Washington 25, D.C. Attn: Technical Library	1
Chief of Research and Development OCS, Department of the Army Washington 25, D.C.	1
Commanding General U.S. Army Materiel Command Washington 25, D.C. Attn: R&D Directorate	1
Commanding General U.S. Army Electronics Command Fort Monmouth, New Jersey Attn: AMSEL-AD	3
Director U.S. Naval Research Laboratory Washington 25, D.C. Attn: Code 2027	1
Commanding Officer Corps of Engineers U.S. Army Map Service Washington 25, D.C. Attn: Library Control No. 62L-525	1
Commanding Officer and Director U.S. Navy Electronics Laboratory San Diego 52, California	1
Commander Wright Air Development Center Attn: WCOSI-3 Wright-Patterson Air Force Base, Ohio	2

Organization	No. of Copies
Commander Aeronautical Systems Division Attn: ASAPRL Wright-Patterson Air Force Base, Ohio	1
Commander Air Force Cambridge Research Laboratories Laurence G. Hanscom Field Bedford, Massachusetts Attn: CRXL-R	1
AFSC Scientific/Technical Liaison Office U.S. Naval Air Development Center Johnsville, Pennsylvania	1
Aeronutronics Ford Road Newport Beach, California Attn: Mr. J. K. Hawkins	1
Chief U.S. Army Security Agency Arlington Hall Station Arlington 12, Virginia	2
Deputy President U.S. Army Security Agency Board Arlington Hall Station Arlington 12, Virginia	1
Stanford Electronics Laboratories Stanford University Stanford, California Attn: Prof. N. Abramson	1
Electronics Systems Laboratory Massachusetts Institute of Technology Cambridge 39, Massachusetts Attn: Lawrence Stark, MD Head, Neurology Section	1
Director National Security Agency Fort George G. Meade, Maryland Attn: CREF-141	1

Organization	No. of Copies
Aeronutronics Fort Road Newport Beach, California Attn: Mr. F. M. Black Acquisitions Librarian	1
Headquarters Electronic Systems Division Laurence G. Hanscom Field Bedford, Massachusetts Attn: ESAT	1
Chrysler Corporation Advanced Projects Organization P.O. Box 1827 Detroit 31, Michigan Attn: Mr. D. N. Buell	1
Commander Air Force Command and Control Development Division Laurence G. Hanscom Field Bedford, Massachusetts Attn: CRZC, Dr. M. R. Nagel	1
Commander Air Force Command and Control Development Division Laurence G. Hanscom Field Bedford, Massachusetts Attn: CCRR	1
Commander Air Force Command and Control Development Division Laurence G. Hanscom Field Bedford, Massachusetts Attn: CCSD	1
Commander Rome Air Development Center Griffiss Air Force Base, New York Attn: RAALD	1
Commanding General U.S. Army Electronics Research and Development Activity Fort Huachuca, Arizona Attn: Technical Library	1

Organization	No. of Copies
Commander Armed Services Technical Information Agency Arlington Hall Station Arlington 12, Virginia Attn: TISIA	10
Secretariat for ETE/JCCE New York University College of Engineering Research Division Building No. 2 401 West 205th Street New York 34, New York	1
Commanding General U.S. Army Materiel Command Building T-17 Gravelly Point, Virginia Attn: AMCRD-RS-PE, Mr. A. Makin	1
IBM Research Center P.O. Box 218 Lamb Estate Yorktown Heights, New York Attn: Mr. A. G. Konheim	1
Minneapolis-Honeywell Regulator Company Ordnance Division 1724 South Mountain Avenue Duarte, California Attn: Mrs. Marjorie M. Ford Technical Library	1
Headquarters Air Force Ballistic Missile Division Air Research and Development Command, USAF Air Force Unit Post Office Los Angeles 45, California Attn: TDC-61-55-8/Mrs. Flory/2364	1
Commanding Officer U.S. Army Electronics Research and Development Laboratory Fort Monmouth, New Jersey Attn: Dr. H. Bennett, SELRA/N-5	1

Organization	No. of Copies
Corps of Engineers Liaison Office U.S. Army Electronics Research and Development Laboratory Fort Monmouth, New Jersey	1
U.S. Navy Electronics Liaison Office U.S. Army Electronics Research and Development Laboratory Fort Monmouth, New Jersey	1
Marine Corps Liaison Office U.S. Army Electronics Research and Development Laboratory Fort Monmouth, New Jersey	1
Commanding Officer U.S. Army Electronics Research and Development Laboratory Fort Monmouth, New Jersey Attn: Logistics Division	3
Commanding Officer U.S. Army Electronics Research and Development Laboratory Fort Monmouth, New Jersey Attn: Director of Research	1
Commanding Officer U.S. Army Electronics Research and Development Laboratory Fort Monmouth, New Jersey Attn: Technical Documents Center	1
Commanding Officer U.S. Army Electronics Research and Development Laboratory Fort Monmouth, New Jersey Attn: Technical Information Division	3
Commanding Officer U.S. Army Electronics Research and Development Laboratory Fort Monmouth, New Jersey Attn: SELRA/NPT	2
Commanding Officer U.S. Army Electronics Materiel Support Agency Fort Monmouth, New Jersey Attn: SELMS-ADJ	1
National Bureau of Standards Washington, D.C. Attn: Dr. S. N. Alexander	1

Organization	No. of Copies
University of Illinois Urbana, Illinois Attn: Dr. Heinz von Foerster	1
Drexel Institute of Technology Philadelphia, Pennsylvania Attn: Dr. R. T. Zern	1
Stanford Electronics Laboratories Stanford University Stanford, California Attn: Prof. B. Widrow	1
Philco Corporation Research Division Blue Bell, Pennsylvania Attn: Mr. S. W. Moulton Associate Director of Research	1
Eugene Pronko Acting Assistant Program Director for Information Retrieval Documentation Research Program Office of Science Information Service National Science Foundation Washington 25, D.C.	1
Radio Corporation of America RCA Laboratories Princeton, New Jersey Attn: Dr. Jack Sklansky	1
Cognitive Systems Research Program Cornell University Hollister Hall Ithaca, New York Attn: Dr. F. Rosenblatt	1
Department of Philosophy University of Michigan 2216 Angell Hall Ann Arbor, Michigan Attn: Dr. A. W. Burks	1

Organization	No. of Copies
Texas Instruments Incorporated Corporate Research and Engineering P.O. Box 5474 Dallas 22, Texas Attn: Technical Reports Service	1
International Business Machines Corporation General Production Division 3605 Highway 52 North Rochester, Minnesota Attn: Olive Smith	1
System Development Corporation 2500 Colorado Avenue Santa Monica, California Attn: Christopher J. Shaw	1
Commanding General U.S. Army Combat Developments Command Fort Belvoir, Virginia Attn: CDCMR-E	1
Commanding Officer U.S. Army Communications-Electronics Combat Developments Agency Fort Huachuca, Arizona	1
Director Fort Monmouth Office U.S. Army Communications-Electronics CD Agency Fort Monmouth, New Jersey	1

This contract is supervised by Mr. William A. Huber, Data Transducer Branch, Data Processing Facilities Division, Communications Department, U.S. Army Electronics Research and Development Laboratory, Fort Monmouth, New Jersey, Telephone Area Code 201-535-1243.

STANFORD RESEARCH INSTITUTE

MENLO PARK CALIFORNIA

Regional Offices and Laboratories

Southern California Laboratories
820 Mission Street
South Pasadena, California

Washington Office
808 17th Street, N.W.
Washington 5, D.C.

New York Office
270 Park Avenue, Room 1770
New York 17, New York

Detroit Office
The Stevens Building
1025 East Maple Road
Birmingham, Michigan

European Office
Pelikanstrasse 37
Zurich 1, Switzerland

Representatives

Honolulu, Hawaii
Finance Factors Building
195 South King Street
Honolulu, Hawaii

London, Ontario, Canada
85 Wychwood Park
London 14, Ontario, Canada

London, England
15 Abbotsbury Close
London W. 14, England

Milan, Italy
Via Macedonio Melloni 40
Milano, Italy

Tokyo, Japan
911 Iino Building
22, 2-chome, Uchisaiwai-cho, Chiyoda-ku
Tokyo, Japan