Conference on

Machine Processing of

Remotely Sensed Data

October 16 - 18, 1973

The Laboratory for Applications of
Remote Sensing

Purdue University
West Lafayette
Indiana

MULTIVARIATE INTERACTIVE DIGITAL ANALYSIS SYSTEM (MIDAS):

A NEW FAST MULTISPECTRAL RECOGNITION SYSTEM*

F. Kriegler, R. Marshall,
S. Lampert, M. Gordon,
C. Connell and R. Kistler

Environmental Research Institute of Michigan (ERIM)**
Ann Arbor, Michigan

ABSTRACT

As a consequence of the quantity and rate of acquiring multispectral data, fast and inexpensive processing is required before practical use may be made of this technique. The speed of this machine is such that an ERTS frame could be processed in forty seconds if the proper tapes were available. Such speeds now make the man operating the system the major limiting factor in performance and requires an experimental matching of the machine to the operator. The system is a prototype, multiple-pipeline digital processor mechanizing the multivariate-Gaussian, maximum-likelihood decision algorithm operating at 2×10^5 pixels/second. It incorporates displays and film printer equipment under control of a general purpose midi-computer and possesses sufficient flexibility that operational versions of the equipment may be subsequently specified as subsets of the system.

INTRODUCTION

A particularly important and easily overlooked aspect of applying a remote sensing, multispectral system to aid in mapping crops, detecting pollution, or locating some ecological disturbance is that of processing the data to provide the proper information to someone in a time short enough to meet his needs. Unfortunately, ongoing programs do overlook this aspect of the system design problem, for reasons which are rarely ever clear.

One would think that, given the utility of these techniques, it would become immediately evident that the data, gathered as it is at very high rates and over very large areas, must be processed before it can be useful. And when the time allowable to produce such results is relatively short, as it is, then it becomes clear that a major problem exists and requires a solution.

The magnitude of the discrepancy between the ability of a sensor to gather this data and the ability of a general purpose computer to process it, then becomes the next aspect of the problem to be assessed, and can probably be best appreciated by considering a brief numerical example. An airborne scanner will, typically, gather data over a 20 to 30 mile flight line in about 15 minutes on one reel of magnetic tape. A general purpose digital computer can classify this data in a time about 1000 times as long as the time required to gather it. Thus, the data collected in 15 minutes will require 15,000 minutes of processing. This would be 6 weeks of processing, working a 40 hour week. Given one such computer to process this data, this aircraft could only be used for eight 15 minute sorties per year. Clearly, the discrepancy in capabilities is unacceptable.

---

*This system is being developed under AAFE Sponsorship (Contract NAS1-11979) and under the technical supervision of W. M. Howle of NASA-Langley.

**Formerly Willow Run Laboratories of The University of Michigan.

At this point, one must examine the problem more carefully. We should point out that the above discrepancy of 1000 to 1 is actually rather conservative, possibly by as much as a factor of 10. We should also point out that this form of processing yields a map containing the color-coded identification of each element of the scene at the time of the overflight, but also containing errors in this color-coding and geometrical errors due to the motions of the airborne sensor. And it is important to note that the scene is not at all unchanging as one flies along; ground conditions vary, the atmosphere varies, the sensor stability varies, and, in short, there is a spatial temporal uncertainty about the data. To be brief, processing, as it becomes better understood, will require even more computation to remove these errors and to improve the results finally provided. A current rationale for this is given by Erickson (Erickson, 1972).

From another point of view, the objective of providing greater speed also provides much lower cost. This is true directly for the operational situation but also for research and development. Given limited resources for research, more investigations can be conducted. It has been estimated that processing costs in an operational system can be reduced by about a factor of twenty or more from some present day feasibility processing costs. For example, the typical data set described earlier could be processed at a cost of less than $400 instead of present day typical costs of about $8000.

Another important objective is the facilitation of control by the operator of the processor. Classification of remotely sensed data is an interactive process in which the man and machine must, in fact, be considered as the real processing system. This is not apparent in some processing systems since the machine is so slow that an operator is easily able to keep pace with the task. In this system, where the processor is substantially faster, the time required by an operator is three or four times longer than that used in processing. Increases in processor speed will then provide little improvement in throughput. It becomes evident that well designed, interactive display and control subsystems will, in reality, offer the greatest gains in throughput. It should also then be evident that the development process must inherently involve an evolutionary refinement of the display and control interface. This is an experimental task, to some extent, and is addressed primarily in the second phase of this development program. It is a task which can be bypassed only at the risk of seriously impaired throughput. This also implies that the system should be considered as a prototype from which further defined systems may be derived for specific programs.

One final motivation exists in addressing this problem. Proper configuration of an operating system almost requires experience in using a system which is similar to that required. It also requires lead times of a year or two to obtain such a system.

One must, therefore, see this problem as one of great and, indeed, critical importance to the objective of making remote sensing a useful, cost-effective tool. It is thus necessary to pursue a course which in the development of these methods, provides assurance that processing methods are available, or can be developed in a timely manner, to meet the demands of those who would use them.

ERIM has been aware of these problems since the initial design and operation of the M-7 scanner and processing equipment at ERIM ten years ago. A real-time processor (SPARC-SPectral Analysis and Recognition Computer) was completed in 1967, allowing analog decision operations using a multivariate maximum likelihood algorithm thus making possible, for the first time, rapid processing of 12-band data. It was then evident that a considerable amount of assistance was needed by the operator in setting up such a machine and controlling its operation. This led to the evolution of a hybrid system employing a general-purpose digital computer to accomplish these functions.

During the period from 1967 to the present time, this equipment (SPARC) has been used to demonstrate the feasibility of many applications of multispectral techniques to problems in hydrology, agriculture, forestry, land use planning, water-pollution, geology, highway construction and in many other fields.

With the motivation that a better processing system will materially assist in practical, economic realization of the benefits of remote sensing, the present system has been conceived and is now being constructed.

APPROACH

A special purpose computer used for recognition (classification) of remotely sensed objects may take several forms. In general, the system design indicates that a general purpose computer be used for statistical analyses and for control and monitoring of the special purpose machine. The special purpose machine could be implemented using analog or digital techniques.

Analog techniques for classifying objects using multispectral data have been well developed over the past few years and, until very recently, appeared superior in cost and complexity to digital techniques in which a special purpose mechanization would be used. One of the most significant cost-components, the multiplier, was previously estimated to cost on the order of several hundred dollars, and would be replicated, in the worst case, about 1000 times. Two significant changes have taken place: 1) the cost per 8-bit multiplier has declined substantially and 2) the speed of such units has increased to a range of 70 to 200 nanoseconds/product. As a result, the multiplier can be built to be time-shared in computing several products per input sample (5 microseconds) so that a factor of cost reduction may be had ranging between 5 and 10. Labor costs for wiring the medium scale integration ALU's (Arithmetic Logic Units) rather than the small scale integration dual in-line packages (DIPS) can be reduced by significant factors also. The classifier in this configuration, then, will have the same organization as the previously designed analog system (SPARC) but will have its functional components implemented in time-shared digital circuitry. Each component will act as a computing element in a hard-wired sequence, which may be thought of as a pipeline or cascade of computational circuits. The desired processing rate will be comparable with the SPARC system, that is, multispectral data collected by aircraft can be processed at a rate which it was collected, i.e., 2×10^5 resolution elements per second.

An ERTS frame (8×10^6 resolution elements) could be classified, element by element, in about 40 seconds provided that the data is recorded in a form such that it could be supplied at a rate of 2×10^5 pixels/second. This is, of course, not the case for CCT's (Computer Compatible Tapes) presently supplied which would allow rates of one-seventh to one-twentieth of that possible with high density tape.

SPECIAL PURPOSE PROCESSING SYSTEM

The ERIM special purpose processing system consists of a conventional digital computer configuration and a special purpose digital processor used for classification of multispectral data. The overall system block diagram is shown in Figure 1.

The general purpose machine is a disk-based Digital Equipment Corporation PDP-11/45 configured with 24K of core, three tapes, serial printer and keyboard-CRT. The system software will be configured about the disk operating system to provide fast access to the operating programs and available languages.

The DEC PDP-11/45 computer system is a fast, medium-size minicomputer, whose architecture allows extremely flexible interfacing with non-standard I/O devices. This ease of interfacing and speed were the primary reasons for use of the PDP-11/45 computer as a controller for the input and classifier hardware used in multispectral data recognition processing.

The hardware peripherals available in the ERIM configuration are, in detail:

1) The PDP-11/45 central processing unit (CPU), including the floating-point processor.

2) An RK-11 disk controller and an RK-05 disk with interchangable cartridges, each containing up to 1.2 million 16-bit words. Average total access time is 90 milliseconds. Data transfer rate is 11.1 microseconds per 16-bit word.

3) 24K of 16-bit core memory with a cycle time of 850 ns, access at 350 ns (450 ns at the UNIBUS).

4) An LA-30 DECwriter data terminal, with a character set of 64 symbols at speeds up to 30 cps. Output is generated by a 5 x 7 matrix.

5) A VT-05 alphanumeric display terminal with a CRT display and communications hardware capable of data transmission at rates up to 300 Baud in full or half duplex modes.

6) A TM-11 magnetic tape drive controller, two 9-track TU-10 magnetic tape drives, and one 7-track TU-10 magnetic tape drive (read/write speed of 45 ips). Densities available are 200, 556, and 800 BPI for the 7-track and 800 BPI only for the 9-track drive.

7) Two (2) DR-11B direct memory access devices for transmission of data between an external device and memory via the UNIBUS, without a need for continuous control by the CPU.

8) One (1) DR-11C device interface for transfer of data between a user device and memory via the UNIBUS.

9) One (1) KW-11P programmable real-time clock, provided programmed real-time interrupts and interval counting in several modes of operation.

Figure 2 shows the configuration in use at ERIM.

Software presently available on the PDP-11/45 system is the Disk Operating System (DOS), Version 8.0, provided by DEC. Under this operating system, MACRO-11 (a Macro assembler), FORTRAN IV, EDIT-11 (a file editing package), ODT-11R (an on-line debugging package), PIP (a file modification and transfer package), and several other utility packages are provided.

All output will be either through the LA-30 DECwriter or will be put on magnetic tape and printed with a line printer. Input is to be done via magnetic tape for large quantities of data or through the keyboard devices.

CLASSIFIER COMPUTATION

The computation which this system performs is a maximum-likelihood decision, assuming a multimodal Gaussian multivariate distribution. This assumption has been well justified at this time by over 100 experiments using multispectral data at ERIM (Marshall, 1971), by a similar number at LARS and, as time goes on, by more and more experience at NASA and other agencies. Although simpler algorithms can perform well for some data sets, a significant percentage of applications demand this powerful decision rule. No penalty in speed and only a small additional cost occurs for using this algorithm, hence it is employed.

Thus, the basic calculation to be performed is

$$\ln\{pr(X)\}$$

where X is the input data vector and the probability density function is a Gaussian density function:

$$\ln\{pr(X)\} = -\frac{1}{2}\left\{(X-M)^T\theta^{-1}(X-M) + \ln|D| + n\ln\left(\frac{\pi}{2}\right)\right\}$$

The

$$(X-M)^T\theta^{-1}(X-M) = Q \qquad (1)$$

is a quadratic calculation in which M is the mean vector for each distribution and θ^{-1} is the inverse of the variance-covariance matrix. This computation must be performed for each object class included in the classification process and the number of computational steps increases as the square of the number of channels.

The calculation given by Equation 1 can be expressed in a number of ways in order to perform the computation. Since the number of bits in the special purpose classifier is limited, it is desirable to express the quadratic calculation in a form in which the calculated number has a very limited range. The variance-covariance matrix θ can be expressed as

$$[\theta] = [\sigma][\rho][\sigma] \qquad (2)$$

where $[\sigma]$ is a diagonal matrix of the standard deviation and $[\rho]$ is the correlation matrix with all 1's on the diagonal and values of 0 to 1 off the diagonal (in some cases negative values may occur). Taking the inverse of (2) yields

$$[\theta]^{-1} = \left[\frac{1}{\sigma}\right] [\rho]^{-1} \left[\frac{1}{\sigma}\right] \qquad (3)$$

Substitution of Equation 3 into (1) results in

$$\left[\frac{X-M}{\sigma}\right]^T [\rho]^{-1} \left[\frac{X-M}{\sigma}\right] \qquad (4)$$

The terms $(X-M)/\sigma$ can have a very wide range. However, if the range

$$-1 \leq (X_i - M_i)/8\,\sigma_i \leq 1 \tag{5}$$

is exceeded the value of X for that channel is too far from the mean to be considered for classification. Truncation of significant bits will occur and a flag is set indicating this condition. This indication is used to reject a decision that the sample is from the particular class.

The computation of Equation 4 could proceed in a straightforward manner, but can be simplified somewhat due to the symmetry of the correlation matrix and its inverse. This simplification can be accomplished in more than one way. One method is as follows:

$$[Y]^T[Y] = \left[\frac{X-M}{\sigma}\right]^T [B]^T[B] \left[\frac{X-M}{\sigma}\right] \tag{6}$$

where B is an upper triangular matrix formed by the decomposition of the inverse rho matrix. By calculating

$$[Y] = [B]\left[\frac{X-M}{\sigma}\right] \tag{7}$$

the final matrix operation is simply

$$[Y]^T[Y] = \sum_{i=1}^{n} y_i^2 \tag{8}$$

where the y_i are the elements of the [Y] vector. The processor is hardwired to perform the above calculations. A block diagram of the processor is shown in Figure 3.

There are four steps implied by Equations 6, 7 and 8. These are:

1) Subtract the mean from each channel.

2) Multiply each result by $1/\sigma$.

3) Perform the Y matrix multiplication on each result of Step 2 to get Y's.

4) Square each resulting Y and add the results together.

Two additional steps complete the classification process:

1) Add the $\ln|D|$ term and

2) Compare the exponents for the smallest value and output its code.

The actual computations are performed by a set of time-shared arithmetic units arranged in a sequence allowing a set of operations whose execution is less than 5 microseconds before the outputs are latched. Each stage supplies its results to a subsequent stage for further processing. Precision varies between 8 and 16 bits as the data progresses through the pipeline and acquires greater significance.

The machine will accept 8 input signals and can classify the results into one of 9 classes (including the null class) at its output. It is expandable to 16 signals and 17 classes. The decision rate is 200×10^3/second.

Internal arithmetic operations are conservatively derated to less than 50% of rated component speeds. For example, an 8 x 8 bit multiply is allowed 300×10^{-9} sec, but could perform at less than 200×10^{-9} sec. As a result, the system should be able to operate at higher speeds later. The general characteristics of the system are summarized in Figures 3 and 4.

Finally, additional equipment is planned for the second phase of this program. This will include a pre-processor to remove spatio-temporal variations of the data prior to classification and improved display equipment to facilitate operator interaction with the system. These equipments are necessary to obtain an effective operating system capable of handling data as it is obtained from the various sensor equipments and to enable effective operator control of the system.

DETAILED OPERATING MODES

The design of the operating modes of the system is of great importance in meeting a primary requirement of the system: fast operation. This is true because when the capability of fast machine processing is available, the time spent by an operator interaction with the system far outweighs the time spent by the machine. This is conservatively estimated to be a factor of 3 or 4 to one. Thus, for 15 minutes of data processing, the operator will spend about an hour in set-up, locating objects and testing performance using an optimized system. This is, of course, much better than is now typical, where the operator may spend 4 to 8 hours in these activities. Thus improvements in raw computing power alone show negligible gains in overall throughput. The greatest gains may be had through facilitation of the operator's task and this cannot be completely accomplished a priori.

The operating modes identified in this system are six:

1) Data Input - in which information is gathered or supplied to specify training and test areas.

2) Analysis - wherein the training set data is processed to obtain parameters for recognition.

3) Set-Up - in which mode the machine is prepared for recognition or enhancement of objects.

4) Test - a mode in which examination of subsets of the data, such as training and test sets, are processed to verify machine setup.

5) Run - the principal machine mode, in which large quantities of data are processed into tabular or pictorial data for output.

6) Diagnostic - a mode enabling isolation of machine faults.

Each of these modes is briefly described below.

DATA INPUT MODE

In this mode, data may be entered from an analog or digital source containing up to 16 analog channels or any number of digital channels at rates determined primarily by the computer compatible tape (CCT) available for output or by the disk. A high density digital tape capability would obviously allow higher speed. The areas entered from the source may be specified by scanner line number, and location along the scan line. Up to sixteen segments of a scan line may be isolated for entry from analog tapes, eight of these from an inertially stable reference and eight from a scanner reference. These segments may be dynamically modified from scan line to scan line in less than 100 microseconds. Data entered is indicated by boundaries overlaid on the scene on a cathode ray oscilloscope (CRO) display (C scope) and by gate pulses for an A-scope display. The basic capability thus exists to allow input along one scan line of at least 8 non-overlapping training or test areas with random boundaries.

The complete data set may also be entered onto disk or tape. Data may be converted from analog to digital, digital to analog (16 channels) 7 track (200, 556, 800 BPI) to nine track CCT (800 BPI) or vice versa.

ANALYSIS MODE

A set of programs will be available to the operator in this mode to allow the calculation of statistical parameters, measurement of scan-angle related radiance functions, calculation of correction functions, calculation of classifier parameters, selection of scanner channels and output to printer and display. These programs have been developed and used, for the most part, on the digital machines (CDC-1604, IBM-7094) at ERIM and have been re-coded for the PDP-11/45.

Additional display-oriented programs will be developed for two-dimensional examination of data under various transformations during the second phase of the system development.

SET-UP MODE

In this mode the parameters for the processing operation are supplied to the classifier by a process of sequentially supplying address and data to each of the random-access memories (RAM's) in the various arithmetic modules of the system. Although this would be a sequential operation during initialization of the classifier, each element of the system set-up sequence contains address and parameter and can access any RAM location. Thus any parameter can be changed, once the change is known, in a few microseconds. For example, the complete processor configuration could be changed in about a millisecond or, again, the set of all mean values could be changed in about 100 microseconds.

The dynamic control of the machine parameters offers a great deal of flexibility in the uses of the machine. For example, adaptive recognition procedures may be implemented given algorithms able to pace the recognition process. Again, the input data may be selected in geometrically varying patterns over the scanner scene. Or, if there were to be a need for such an operation, the set of inputs selected could be changed in a few microseconds.

Set-up also includes the selection of display data and the display format. The data access is under RAM control and may be modified by changing RAM parameters for the display bus selection.

TEST MODE

Once a set-up has been made, the usual operation is to process a large quantity of data to obtain a recognition map or a tabulation of data. Since this may take a few minutes and much paper or film, it is desirable to obtain a sample of results obtained by processing the training and test sets stored in the machine to determine that the procedure has operated properly.

The operator may call for a display of the recognition results over one or more training or test fields. Problems due to variations of data over a field may be observed and corrected, if present. Tabulations of results for each field may be made and presented as percentages of various categories for the classes of interest.

Essentially, this mode is one of examining performance of the machine using data already stored in memory as a part of the input process. The machine is thus classifying digital data which, in this case, is training and test data but could also be all the data the operator may wish to process. The test mode in this case may be considered the final step in processing.

RUN MODE

In this mode, the input tape is run over the data gathered from the scene of interest and the data is supplied to the processor directly from a digital source or via the 16 A-D converters. Several modes of processing can be used. All the data may be processed and the classified output may be employed to create a recognition map or may be tabulated to obtain the number of acres of various objects in the scene. Again, (making use of the dynamic control of scan line segmentation) all the data may be processed and results tabulated on a per-field basis by sorting the results into bins for each field determined by known boundary locations.

Data classified as to type may be displayed as it is processed on a C-scope with a long-persistence phosphor so that the operator in real time can verify the performance of the system. In the second phase of the program, this C-scope will be replaced with a color moving window display able to show all objects classified at one time as though the operator were looking through a window in the bottom of the scanner aircraft and were able to see each scene point in a particular color corresponding to its classification. The data may also be recorded in color on film for closer examination at higher resolution.

DIAGNOSTIC MODE

Since the processor contains a large number of circuits and is reasonably complex, it has been designed to allow access to outputs of most major components in the system to reduce down time for fault correction. The machine operation may be verified by supplying known inputs and ascertaining that the proper codes are observed at the output and intermediate points. For this purpose, a diagnostic/display bus system is incorporated and can address the desired points for input to the general purpose machine or for display on a CRO, the moving window display or on film.

CLASSIFIER OPERATION

The classifier consists of a number of "pipeline" computers in parallel as shown in Figure 3, finally converging on the circuits to scale the exponents of the density function and intercompare these exponents for a decision. The sequence of operations can be visualized as shown in Figure 6 where data is shown entering the A-D converters at t = 0 in the upper left corner of the diagram. A sample, consisting of a vector of eight elements of eight bits each is passed through the computational circuits indicated and emerges at the bottom right of the diagram as a classification code of 5 bits. The general appearance and function of the arithmetic operations so diagrammed is that of a "cascade" in which the breadth of the cascade in time is proportional to the computational load of a particular circuit.

Each "pipe" or "cascade" processes the computation of the quadratic form for the exponent of the Gaussian distributions for two distributions. There are, then, a sequence of alternating computations of the first exponent in cascade 1, the second exponent in cascade 1, the third exponent in cascade 2, etc. In the machine, at present, there are four such cascades operating in parallel, allowing the computation of eight exponents at once. The machine will be expanded to eight cascades to allow computation of 16 distribution exponents at a time.

The timing diagram, Figure 6, shows the flow of two exponent computations and the resulting decision, neglecting the fact, for the sake of simplicity, that the cascade would normally contain portions of other computations for the preceding and subsequent samples. Data is entered into the A-D converters and is available for computation at the end of 16 machine cycles, approximately 5 microseconds. Data latched in the converter outputs is then supplied sequentially (X_1, X_2,...X_8) via a multiplexer and subtracter to a latch, and one clock cycle per element is allowed for this operation. At $t = 1$, μ_1 is subtracted from X_1; at $t = 2$, μ_2 is subtracted from X_2, etc. Next, at $t = 2$, $(X_1 - \mu_1)$ is multiplied by $(1/8\sigma_1)$ yielding X_1^*. At $t = 3$, the result is supplied to each of 8 multiplier-summers which compute the products $(X_1^* \cdot P_{11})$, $(X_1^* \cdot P_{21})$, ...$(X_1^* \cdot P_{81})$ and enter these into the summers. At $t = 4$, the multiplier-summers compute $(X_2^* \cdot P_{12})$, $(X_2^* \cdot P_{22})$.... · ....$(X_2^* \cdot P_{82})$ and add these products to the previous results. Thus at $t = 11$ the summers contain the complete sums of products for all matrix operations. Each of the eight multiplier-summers may, as a result, be considered as a row operator since it accomplishes the sequence of multiplications and summations for a particular row.

This vector, the result of a transform which makes the product pairwise uncorrelated, need only have its elements squared and summed to obtain the normalized quadratic form for the exponent. This is accomplished during cycles $t = 11$ to $t = 19$, allowing one cycle for each squaring operation and a final cycle for storage of the summation.

At this point, the exponent must be re-scaled and the natural logarithm of the determinant of the covariance matrix added to obtain the final exponent of the density function. This requires two cycles. The comparison of these exponents, now becoming available from the normalization circuitry, begins as each exponent appears. The procedure is to examine all exponents sequentially to choose the smallest, assuming that one is less than a threshold test value which is entered first, and to retain at all times the lesser value of two sequentially examined exponents. The number of the exponent retained specifies the class of the input vector. This is available to be displayed, printed on film or supplied to the computer for logging or subsequent processing.

The cascades may also be used to process an increased number of distributions for a lesser number of channels. Thus for a 4-channel source, such as ERTS, the operations of the various arithmetic units may be time shared to provide 16 class decisions instead of 8 class decisions for 8 channels.

CONCLUSION

The MIDAS system will provide a significant increase in the capability to process the ever-increasing load of multispectral data. It will also make research investigations requiring substantial amounts of computation much less costly than at present. Finally, MIDAS will provide needed experience in operating a high-throughput system as a predecessor for operational systems.

REFERENCES

Erickson, J. D., "Automatic Extraction of Information from Multispectral Scanner Data", International Archives of Photogrammetry, Volume 17, January 1973.

Marshall, R. E. and Kriegler, F. J., "An Operational Multispectral Surveys System", Proceedings of the Seventh International Symposium on Remote Sensing of Environment, Ann Arbor, Michigan, May 1971, p. 2169.

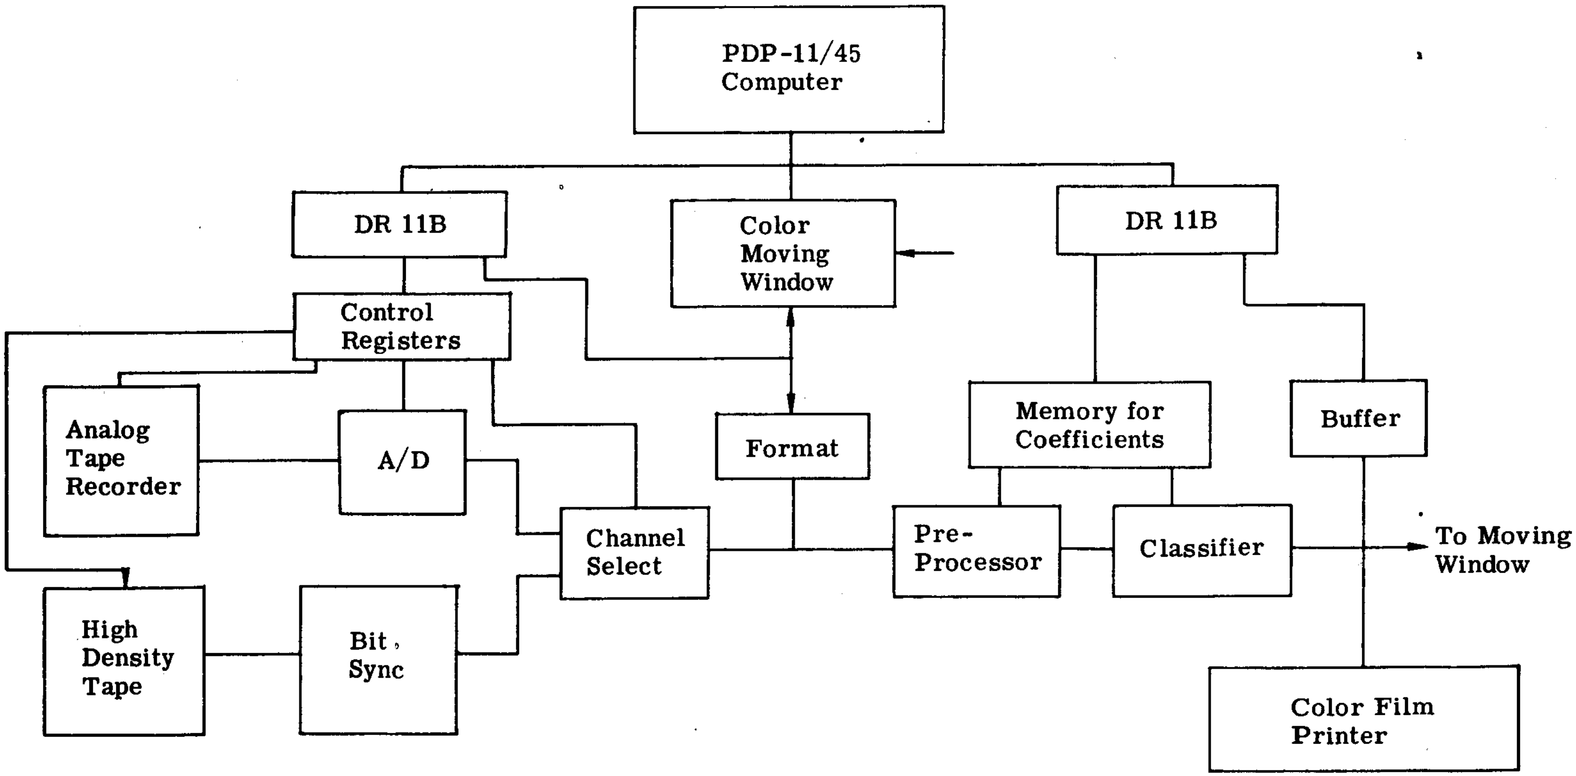

Overall System Block Diagram

Figure 1

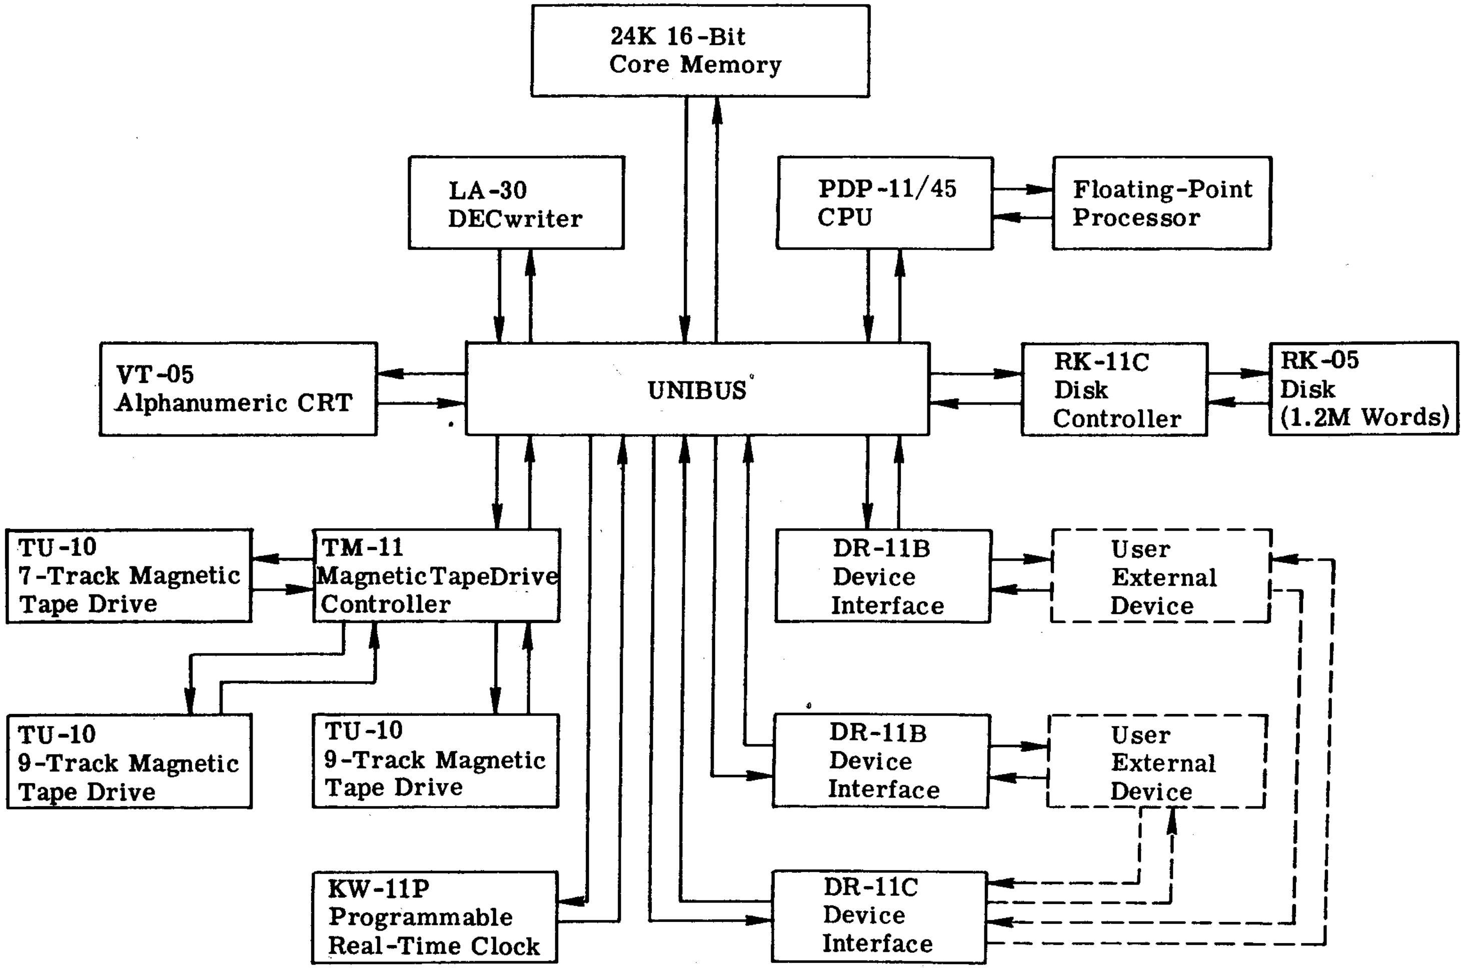

PDP-11/45
CONFIGURATION
Figure 2

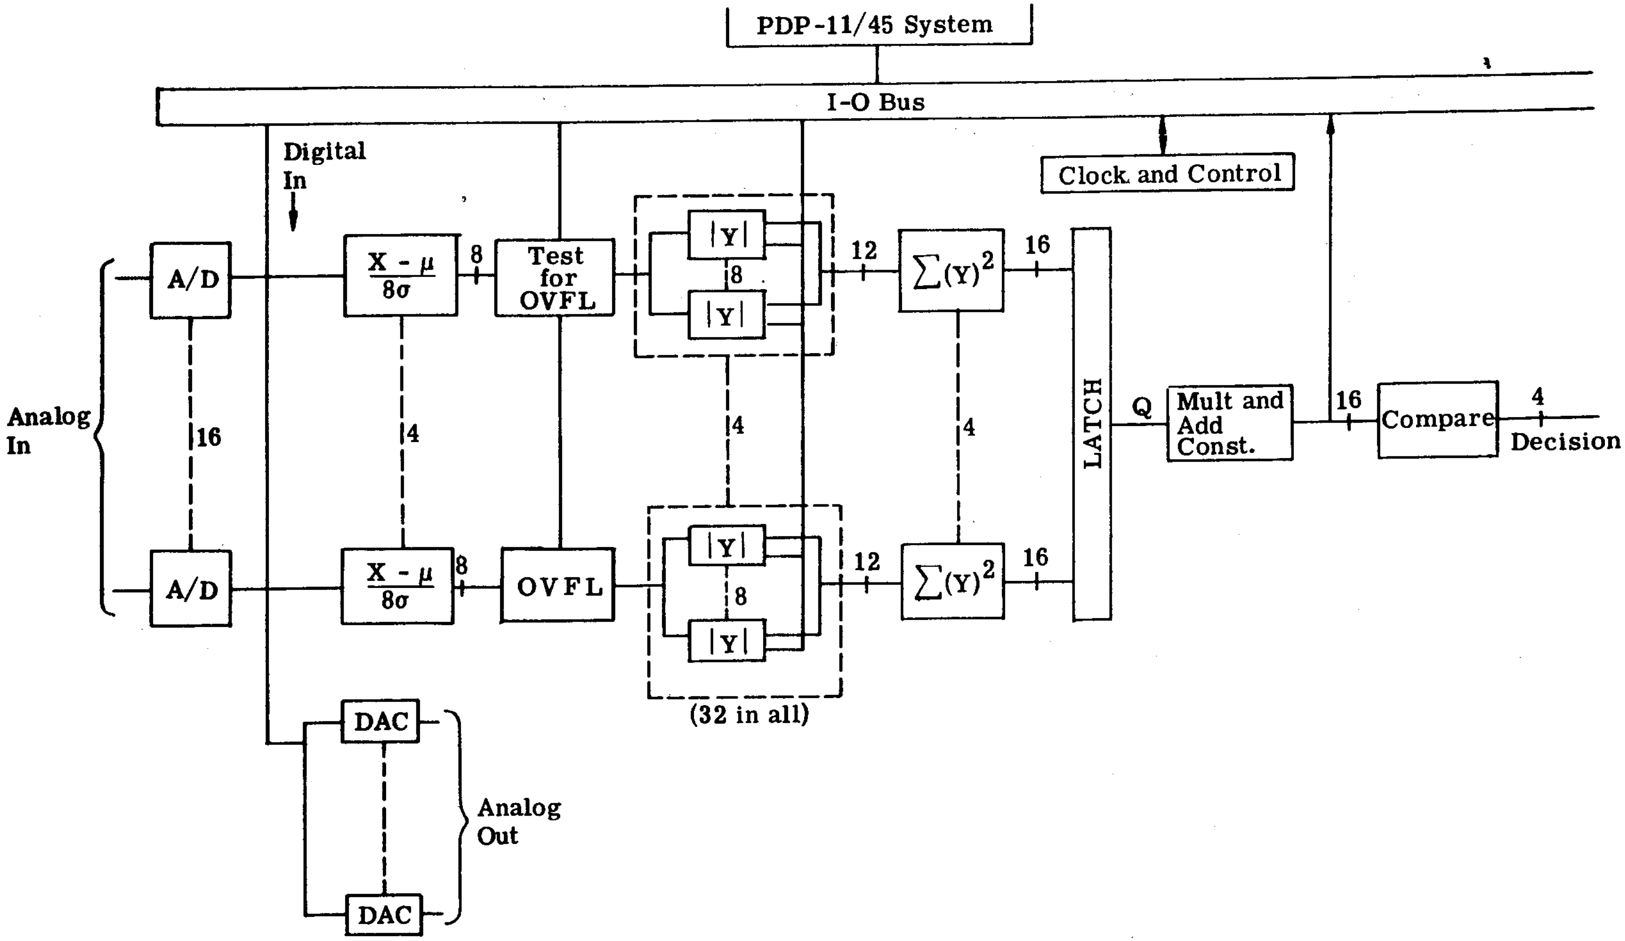

BLOCK DIAGRAM OF PARALLEL PROCESSOR

Figure 3

ADVANTAGES OF DIGITAL PARALLEL CLASSIFIER

Uses Current State-of-the-Art Digital Techniques

Less Costly Than Current Hybrid Techniques

Complete Repeatability in Set-Up and Performance

Computer Controlled Diagnostics Easily Implemented for Error Free Operation

Throughput Equal to Current Analog/Hybrid Techniques

Figure 4

CHARACTERISTICS OF PROTOTYPE PROCESSOR

All Digital Parallel Classifier Under Computer Control

Classification Rate of 200,000 Data Vectors/Second

Classification of Data Vector Into One-of-Eight Stored Categories for Eight Channel Data

Classification of Data Vector Into One-of-Sixteen Stored Categories for Four Channel Data

Potential of Expansion to Sixteen Channel Capability

Figure 5

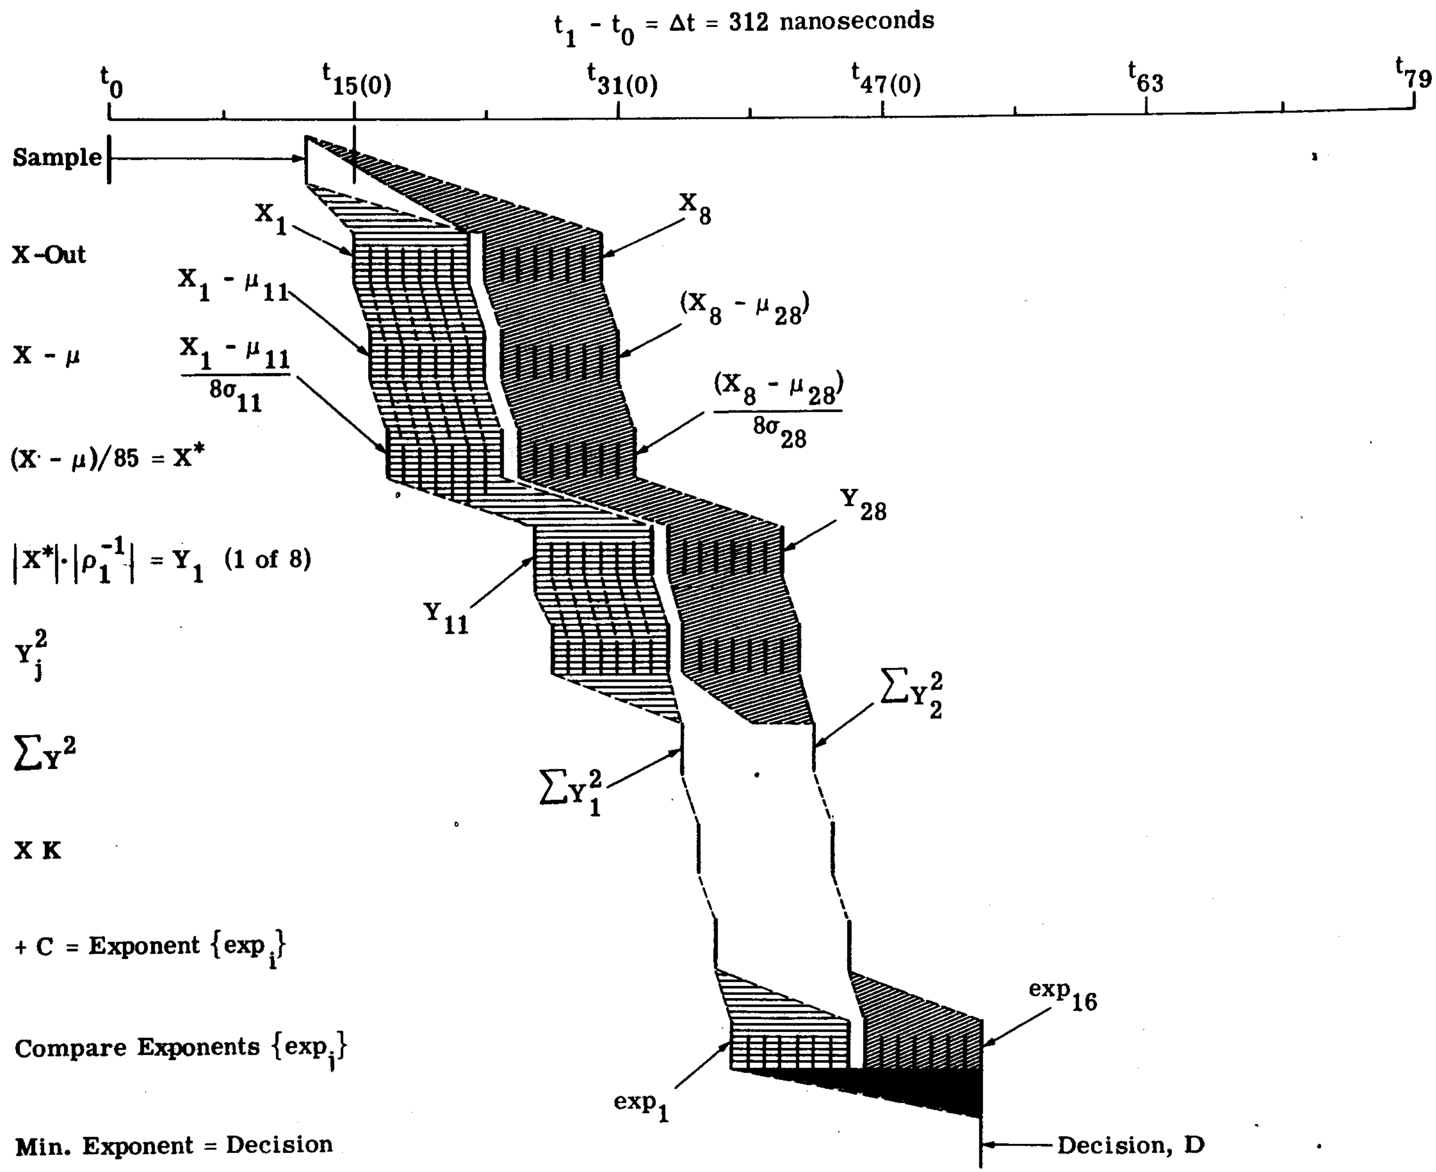

Figure 6. **CLASSIFIER TIMING DIAGRAM**