

Final Report — Phase II

January 1971

Covering the Period 18 January 1968 to 1 February 1971

RESEARCH STUDY OF A FUNDUS TRACKER

By: D. H. KELLY H. D. CRANE

Prepared for:

NATIONAL AERONAUTICS AND SPACE ADMINISTRATION
AMES RESEARCH CENTER
MOFFETT FIELD, CALIFORNIA 94035

CONTRACT NAS2-3995


SRI Project 6319

Approved by:

F. J. KAMPHOEFNER, *Director*
Engineering Sciences Laboratory

B. COX, *Executive Director*
Information Science and Engineering Division

Copy No.



ABSTRACT

With minor exceptions, previous stabilized-image experiments with the human eye have used contact-lens apparatus, which has various disadvantages. Cornsweet proposed a stabilization technique based on tracking the edge of a blood vessel in the optic disk, which we have developed as follows. A small spot of blue light is imaged on the fundus, and scanned in a circle around the optic disk at high speed. Variations in reflectance along this circular raster are detected by a photomultiplier tube and digitized for further processing. A selected frame of this digital video signal is stored in a magnetic-core array and correlated with subsequent frames to generate eye-movement information in real time. The correlation algorithm distinguishes horizontal, vertical, and torsional motions of the fundus pattern. Analog displacement signals fed back to a servo-mechanism can thus cause the raster to track the fundus pattern, locking to the position of the selected frame. Such closed-loop signals can be used simultaneously to move a stimulus pattern in stabilized-image experiments.

CONTENTS

FRONTISPIECE

ABSTRACT	ii
LIST OF ILLUSTRATIONS	iv
I INTRODUCTION	1
A. History	1
B. Principles of Operation	3
C. Model Eye Development	7
D. Other Optical and Mechanical Improvements	9
II COMPUTER SIMULATION WITH CRT	11
A. Slow-speed Scanning	12
B. Signal/Noise Extrapolations	12
III NEW SCANNING SYSTEM	15
A. Optics	15
B. Sync Generation	17
IV COMPUTER SIMULATION WITH NEW SYSTEM	19
A. Video Display Technique	20
B. Tracking Records	21
V DESIGN AND CONSTRUCTION OF SIGNAL PROCESSOR	26
A. Scanning Rate and Sampling Rate	26
B. Overall Timing	28
C. Forming the Master Frame	29
D. Tracking Mode	31
E. Align Mode	33
VI PRESENT STATUS	35
BIBLIOGRAPHY	38

ILLUSTRATIONS

Figure 1	Subsystems of the Fundus Tracker	3
Figure 2	Algorithm Used for Signal Processing	5
Figure 3	Waveforms in the Cross-correlation Process	6
Figure 4	Cross-section of Model Eye	8
Figure 5	Direct Computer Simulation at Low Speed	11
Figure 6	Trade-offs in Tracking Accuracy	14
Figure 7	Configuration of the Mechanical Scanner	16
Figure 8	Computer Simulation with High-Speed Recording	19
Figure 9	Video Signal from Human Fundus	20
Figure 10	Video Signal and Computed Position Outputs for Voluntary Eye Movements	22
Figure 11	Trajectory of Eye-movements Constructed from Figure 10	23
Figure 12	Filtered Video Signals and Computed Position Outputs for Involuntary Eye-movements	24
Figure 13	Overall Schematic Diagram of Signal Processor	27

I INTRODUCTION

This report describes a new method of tracking small eye-movements and stabilizing the retinal image. The importance of stabilized-image experiments in vision is well established. The only successful methods of image stabilization to date involve the attachment of a tightly-fitting contact lens to the subject's eye, but this has certain disadvantages. Usually a mirror attached to the contact lens becomes part of an optical system that compensates for eye movements. Although this technique can probably eliminate slippage with respect to the cornea, there remains a short-term stabilization error of about half a minute of arc.

Since the eyeball is not a rigid structure, perhaps a significant fraction of this error is attributable to motions of the fundus with respect to the cornea, which cannot be eliminated by the contact-lens technique. Discomfort and risk of ocular damage are also not trivial; subjects frequently report a temporary blurring or distortion of vision after the lens is removed.

We have undertaken to develop an alternative stabilization technique, not involving contact lenses. Our approach is to optically track the motion of the fundus itself, and to use this eye-movement information in a servo system for stabilizing any stimulus pattern with respect to the retina.

A. History

Phase I of this program, which occupied most of the year 1967, was a technical study to determine the feasibility of a new method of tracking

the involuntary movements of the human eye, partly for the purpose of stabilizing the image on the retina. The proposed method involved tracking the fundus itself (where the retina is located), in contrast to various other methods which track the motions of the cornea, lens, or sclera. The conclusion of Phase I was that it would probably be possible, within the present state of the art, to devise a fundus tracker for this purpose, by appropriately choosing the optics, wavelengths, and region of the fundus to be tracked, and with essentially no risk of physiological damage (there is some risk with contact-lens methods). Some tests of prototype components were also conducted. For more details of this feasibility study, the reader is referred to the Final Report on Phase I, dated November 1967 (Contract NAS 2-3995).

In Phase II, which began 18 January 1968, we undertook to design and build a prototype fundus tracker, using the concepts, principles and components evolved in Phase I. By April of 1969 we had constructed a prototype which was almost complete, but it lacked an essential subsystem (known as the digital signal processor) whose function is to convert the retinal-reflectance signals into position-error signals, by means of an extremely efficient noise-suppressing algorithm developed in Phase I. At that time, the development work was interrupted for lack of funding; the signal processor had been designed, but only partly constructed. In June, 1970, work was again resumed, with the goal of completing the signal processor and getting the prototype tracker into operation.

Most aspects of this program have proceeded remarkably well. Its most serious technical setback occurred during the first six months of 1968, when accurate signal/noise data were obtained for the first time, with the aid of a computer program that simulated the operation of the signal processor. These calculations revealed that the light level required for our desired tracking accuracy (although below the threshold

of retinal damage by a factor of more than 10^3) was still about two orders of magnitude greater than the maximum luminous energy available from the CRT source in our original scanner. This necessitated re-designing the system, to convert it to a mechanical scanner and Hg-arc light source. This unforeseen development, and later a series of defective memory-core arrays in the signal processor, caused some delay in the prototype development. An additional year was lost for lack of funding.

The final status of the fundus tracker at the end of Phase II is discussed in Section VI.

B. Principles of Operation

Here we will briefly review the basic principles of the fundus tracker. Its main components are shown in Figure 1. The light source is a flying-spot scanning pattern which is imaged on the fundus near the optic disc. A photomultiplier tube detects the resulting video signal reflected from the fundus pattern. The signal processor derives the




FIGURE 1 SUBSYSTEMS OF THE FUNDUS TRACKER

desired motion signals from the video signal by means of a correlation algorithm described below. The flying spot can then be deflected accordingly by a servomechanism which adds the fundus motion signals to the scanning motion.

The switch S represents the possibility of open- as well as closed-loop operation. Since the outputs of the signal-processing algorithm are linear with fundus displacement over a small range, they can also be used to indicate small eye-movements directly, without closing the loop.

In order to produce a stabilized image, the feedback switch S is closed, so that the scanning pattern is locked onto a fixed location on the fundus. The motion signals then become equal to the actual fundus displacement, and can be used to stabilize other patterns projected on the retina.

As shown in the frontispiece, the scanning pattern is simply a circle about 6° in diameter, positioned near the optic disc to cross as many edges as possible in all directions. The motor rotates at 24,000 rpm, so that the scanning rate is 400 frames per second around this circular path. The scanning circle is 15° or 20° off the visual axis; the fovea is near the left side of the picture, surrounded by small capillaries. The choice of this scanning pattern was based on extensive tests conducted in Phase I.

Figure 2 shows how the signal processor derives motion signals from the video signal. First a master frame is formed and stored in the memory; this master frame consists of the algebraic difference between one video frame and a slightly delayed version of itself. Selection of this frame determines the location on the fundus to which the scanning circle will be locked in the closed-loop condition. The master difference frame is then stored and multiplied repeatedly by each new frame of the video signal as it comes, in real time. This product is integrated over





FIGURE 2 ALGORITHM USED FOR SIGNAL PROCESSING

the duration of one frame, yielding a value which is positive, negative, or zero, depending on whether the incoming video signal was earlier, later, or the same as the master frame (relative to the sync pulse). This calculation corresponds to a pure rotational component about the center of the scanning circle. The X and Y motion signals are obtained in the same way, except that before integrating, the video signal is also multiplied by a gating signal or weighting function, which consists of the sine or cosine of the angular position of the flying-spot. In the actual signal processor described in Section V, we used square waves instead. These are much simpler to generate digitally, though they introduce other problems, which are discussed in Section VI. At present, we feed back only the X and Y signals for image stabilization, but the rotational component is also available if needed.

Figure 3 illustrates this algorithm with an artificial fundus pattern consisting of one broad, vertical blood vessel. The video signal then consists of two broad pulses, 180° apart, as shown in line 1. The master difference frame formed from this input will contain narrower



TB-6319-4

FIGURE 3 WAVEFORMS IN THE CROSS-CORRELATION PROCESS

pulses, positive at the leading edge and negative at the trailing edge of each video pulse. Now at some subsequent time, let the fundus be displaced in the X direction. The incoming frame then is shown in line 5, with the first pulse advanced and the second pulse retarded relative to the master frame. When this is multiplied by the stored master frame, it leaves a positive pulse in the first instance and a negative one in the second, as in line 6. Note that this product alone would integrate to zero (as it should, since we have postulated no rotational movement). But when it is multiplied by the sine θ gating signal in line 7, the polarity of the second pulse is reversed, and the result now integrates to a positive value for X-displacement.

If the X displacement had been in the opposite direction, both pulses in the last line (and hence its average) would have been negative. An entirely analogous derivation of the Y displacement signal is obtained by integrating the product $f_{\delta} f(\theta) \cos\theta$. In the illustrated case of X movement only, this product would of course integrate to zero.

As in any correlator, major noise suppression results from integration over a complete frame; the use of the quadrature components makes this averaging as effective as possible without throwing away any position information.


In practice, these operations are carried out entirely in digital form; i.e., the signal processor is essentially a special-purpose, hard-wired computer (see Section V). At this low light level, the video signal can be obtained in digital form by counting photomultiplier output pulses with a 100-MHz counter. We collect this count 512 times per frame, so that each sampling interval occupies about 5 microseconds at the present scanning rate. The average count is less than 80 events during this interval. All subsequent processing is done with 8-bit precision, so the only significant noise in the open-loop output should be the result of photon noise in the original video signal. These numbers are then processed by a modified version of the correlation algorithm just described, which provides essentially optimum filtering of this noise.

C. Model Eye Development

In testing our design for the fundus tracker, we needed an input that would realistically resemble the fundus pattern of the human eye, but would remain perfectly stationary for hours, if necessary, or could be moved by precisely known amounts. In order to test all parts of the optical and electronic hardware as they became available, it was obvious

that this test pattern should be embedded in a model eye, located in front of the fundus scanner, and readily interchangeable with a real human eye whenever desired.

After determining that simpler devices were not adequate, we constructed the model eye shown in Figure 4. The artificial cornea is



TA-6319-43

FIGURE 4 CROSS-SECTION OF MODEL EYE

actually a large (scleral) contact lens of zero power; this is cemented to a brass cell, which was filled with a liquid to simulate the fluid vitreous. The artificial fundus is actually a photograph of a real human fundus on Microfile film (painted white on the back, to convert it into a reflectance target). With water as the liquid, an unacceptably bright specular reflection was obtained from the artificial fundus, due to the mismatch of refractive index between the water and the film base. This problem was eliminated by changing the liquid to glycerin. The optics of the model eye were thus very similar to those of the real eye

(for which the optics of the fundus scanner were designed), and hence gave excellent image quality. In addition to the realistic fundus pattern, other targets were also used. A radial spoke pattern was found very useful for aligning the system, checking the computer simulation program, and making other measurements.

D. Other Optical and Mechanical Improvements

Early in Phase I, it had been determined that we would not attempt to design and construct our own optical system for the fundus tracker. A much more satisfactory solution (from the viewpoint of better image quality as well as lower cost) was to buy a commercial fundus camera, and replace its light source with our photomultiplier, and its film carrier with our flying spot. At the end of Phase I, the (then) new model of the Zeiss fundus camera had been selected for this purpose, and modified to accept a CRT and ten-stage photomultiplier, also carefully selected for this purpose. (The CRT was later replaced by a mechanical scanner consisting of a mercury arc, optical fiber, lenses, filters, beamsplitters, scanning mirror and high-speed motor, as described in Section III.)

The clinical headrest and adjustment platform which came with the Zeiss camera were too shaky for any precise optical experiments. They were removed and several other improvements made in the interest of rigidity. All of the optics, including a bite-board mounted in a lathe cross-feed arrangement, were bolted to a 4' x 4' plate of 1/2" aluminum. In order to minimize vibration, this plate was supported by 2" of isolating foam material, on a very heavy wooden table. Although slightly less comfortable for the subject, this arrangement has been much more satisfactory for the experiments.

Considerable work was necessary in order to get the special CRT to perform according to the manufacturer's specifications. Several controls were installed for test purposes, to conveniently adjust the size and location of the circular path of the scanning spot. By March of 1968, based on both model-eye and real-eye tests, we felt that we had the best video signal obtainable from such a CRT scanner, and we were ready to proceed with the problem of converting this into an eye-movement signal, suitable for tracking the motions of the retina.

II COMPUTER SIMULATION WITH CRT

The only satisfactory method of estimating the open-loop performance of the final device, in terms of speed and precision of tracking, was to actually carry out the correlation algorithm and study the results of these calculations for various input conditions. A computer program written for this purpose in Phase I was available to test the characteristics of the CRT scanner, as illustrated in Figure 5. Since the signal




FIGURE 5 DIRECT COMPUTER SIMULATION AT LOW SPEED

processor that we later built also performs these calculations on a digital basis, the computer simulation of its performance was free of sampling artifacts. However, the computer program could not run fast enough to operate in real time. (Also, the computer program used sine and cosine factors instead of square waves to calculate the X and Y motion components; see Section VI.)

A. Slow-Speed Scanning

In order to track the fastest eye movements, a desirable scan rate for the final device is between 10^2 and 10^3 scans per second. The latter figure is close to the limit of the state of the art, even for a hard-wired single-purpose device. (We finally settled for 400 scans per second, but this was limited as much by the scanner as by the signal processor.) In our computer simulation tests with the CRT scanner, the CRT deflection circuits were modified to give a very slow scanning rate (5 sec per frame). This was required in order to give the software in the SDS 910 computer time enough to keep up with the input data.

This slow rate also greatly increased the signal/noise ratio, but from the resulting data we were able to calculate the (open-loop) characteristics for other scanning rates, spot sizes, etc., using as our input the realistic model eye described in the previous section. It was not possible to get two closely similar, successive scans from a real eye under these conditions, of course, or to attempt any closed-loop experiments. However, we did compare the video signal from a real eye with that from the model eye, both at a scanning rate of 100 frames per second, in order to assure ourselves that they were not significantly different.

B. Signal/Noise Extrapolations

Since the only significant noise in the system is the quantum noise inherent in the photodetection process, extrapolation of the simulation tests was relatively straightforward. It can be shown that the signal-to-noise ratio (and hence the tracking accuracy) increases with the square root of the amount of light available (i.e., the product of area times brightness of the flying spot) and decreases with the inverse square root of the scanning speed. However, increasing the spot size above the diffraction limit decreases the signal bandwidth so there is

obviously an optimum spot size, and corresponding low-pass filter; these optima depend strongly on the fine detail in the fundus pattern.

These simulation studies showed that no CRT could produce a bright enough flying spot to provide the desired tracking accuracy; therefore we concluded that a different type of flying-spot light source was required. We conducted a number of experiments optimizing the spot size, the radius of the scanning circle, the delay in the cross-correlation algorithm, and other parameters, and from these we estimated the optimum tracking accuracy with the CRT to be 2 or 3 orders of magnitude lower than needed.

Our calculations concerning this and other potential light sources are summarized in Figure 6. As this figure shows, a high pressure mercury arc should permit us to track the fundus with a precision of less than one minute of arc at scanning rates greater than 100 Hz; but at a brightness that is still three orders of magnitude below the damage threshold for the retina. We therefore decided to use a mercury-arc light source in the mechanical scanner described in the next section.





FIGURE 6 TRADE-OFFS IN TRACKING ACCURACY WITH THE SIMULATED FUNDUS

III NEW SCANNING SYSTEM

As a result of the tests described in the previous section, we set about to devise a mechanical scanning system that could use a high-pressure mercury arc as its light source. The new optical system is still based on a modified Zeiss fundus camera, from which the xenon flash-tube had been removed and replaced by a blue-sensitive photomultiplier tube in Phase I. In the present modification, however, we replaced the CRT (which replaced the original film carrier) with our new flying-spot scanner assembly, consisting of a mercury arc, optical fiber, filters, lenses, beam-splitter, scanning mirror and high-speed motor.

A. Optics

As shown in Figure 7, the mirror is mounted on the shaft of the high-speed motor. Light from the mercury arc is focused on one end of the optical fiber; the other end is located in the focal plane of the optical system. Light from this end of the fiber is reflected by the beamsplitter to the spinning mirror, and then transmitted back through the same beamsplitter and eventually focused on the subject's fundus. This arrangement provides the only satisfactory method of accommodating the new scanner to the Zeiss optical system. Absorption losses in the fiber and beamsplitter are negligible, but the reflection-transmission loss due to the beamsplitter limits the efficiency to 25%. The fiber is required for closed-loop operation, as described below; it also performs the necessary function of a pinhole. A 5-mil diameter fiber was chosen for this purpose, on the basis of the tests described in Section IV.



TA-6319-3R

FIGURE 7 CONFIGURATION OF THE MECHANICAL SCANNER

The mercury arc is a 100-watt PEK type 112-2118, continuously cooled by an air jet. A microscope objective images the arc onto the end of the fiber, which is mounted in a 3-way rack-and-pinion arrangement for precise focusing and positioning. A custom-made interference filter at the lamphouse eliminates from the system all wavelengths less than 400 nm or greater than 500 nm. Subjectively, the spot appears to be about the same color as the former CRT spot, but it provides more than 200 times as much light.

This bright blue spot undergoes two types of motion on the subject's fundus: scanning and tracking. The scanning motion is imparted by the spinning mirror, which is mounted at a slight angle with respect to the motor shaft. This moves the image of the end of the fiber in a circular path when the motor is rotated. At 400 rps, it has the appearance of a steady circle about 6° in diameter; normally this circle is imaged at

the subject's optic disk. The tracking motion is provided by deflecting the output end of the fiber itself; this is done by two electromechanical driver units (modified loudspeakers) that are mounted at right angles to each other; these in turn are driven by the X and Y servo outputs of the electronic signal processor (see Section V). Thus, in closed-loop operation, the scanning circle follows the motions of the fundus.

This circle is slightly larger than the optic disk and therefore sometimes visible to the subject, especially in the open-loop condition (scattered light from it would probably be detectable even if it were wholly inside the disk). But since it is located in the periphery of the visual field, it does not interfere with normal visual tasks. When accurately stabilized, of course, it should disappear in any case.

B. Sync Generation

The scanning motor is a 400-cycle synchronous, split-phase type (Globe 18A103) which rotates at 24,000 rpm, providing a frame rate that is fast enough for tracking most eye-movements and compatible with the signal-processing equipment (see Section V). Toothed wheels are mounted on the motor shaft for sync generation, with fixed magnetic pickup heads.

The way in which the mechanical scanner synchronizes the operations of the digital signal-processor is one of its key features. For this purpose, sync-pulses are provided not only for every revolution of the motor shaft (frame sync) but also for every picture element scanned (minor sync). This makes the precision of processing the signal essentially independent of variations in scanning rate, which would otherwise decrease the tracking accuracy.

The 512 minor sync pulses per revolution are obtained by frequency-doubling the output of a magnetic pickup which senses the edge of a wheel containing 256 teeth. Another wheel with one tooth and another

pickup provide the major sync pulse. In this way, the signal processor is effectively locked to the desired positions in space (not to fixed intervals of time), so that slight variations in the speed of the motor affect only the gain, not the waveform of the digitized video signal.

IV COMPUTER SIMULATION WITH NEW SYSTEM

As soon as the mechanical scanner was operating, we tested its open-loop performance, again by computer simulation. That is, we programmed a general-purpose computer to simulate the functions of the signal processor. We slowed the scanner down to run at only 50 frames per second, but even this scanning rate was too fast for the software processor to handle the video signals in real time. We therefore recorded the video and sync signals on magnetic tape, spliced the tape into a loop and played it back at a much slower speed to the computer, as shown in Figure 8. In this way, we obtained real eye-movement records (but not in real time) before the signal-processing equipment was built.




FIGURE 8 COMPUTER SIMULATION WITH HIGH-SPEED RECORDING

A. Video Display Technique

Since the playback of these video recordings into the simulator program was very time-consuming, we had to select only the most suitable segments for eye-movement computation. This was done by displaying the video signal as an intensity-modulated line-scan on an oscilloscope and photographing each frame in succession with a continuously-running Grass camera. The result is shown in Figure 9.




FIGURE 9 VIDEO SIGNAL FROM HUMAN FUNDUS

Here each vertical line represents one circular scan at a rate of 50 scans per second. Since polarity is preserved, the dark streaks are where the flying spot crosses a blood vessel and the lightest ones are probably parts of the optic disk. If there were no eye-movements, these streaks would of course all be straight, horizontal lines. By inspection, most of the eye-movements in these video records appear to be displacements rather than rotations. In other words, no one frame is uniformly advanced or retarded with respect to the next; whenever there is a jump in one horizontal streak, there seems to be an equal and opposite jump at

a point 180° around the scanning circle. On the basis of many video recordings like this, we chose a few segments for computer processing.


B. Tracking Records

Figure 10 shows a computed record of very small but voluntary eye-movements guided by four fixation points, 20 minutes apart in a square array; the subject changed his fixation from one corner of this square to the next as rapidly as he could. The open-loop X-output is shown at the top of the video recording and the Y-output at the bottom. Note the correlation between the X and Y displacements. In the open-loop condition, this can occur as an artifact if the input signal power is not distributed uniformly among the four quadrants of the scanning circle. However, inspection of the video signal shows the effect is mainly real in this case. That is, the subject seldom makes a pure X or pure Y movement under these conditions.

This is more evident in Figure 11, where we have dropped the time parameter and plotted the trajectory corresponding to the X and Y displacements of the previous figure. Here we can see the subject's fixation jitter around one point for a fraction of a second and then leap to another one. The path is not square or rectangular but an irregular tetragon; it does not even return to the same starting point. This behavior is typical of small, rapid, voluntary eye-movements with no moving target to pursue.


Finally, Figure 12 shows a record with no voluntary movements but several involuntary saccades during the course of a few seconds. The saccades vary from 5 to 10 minutes of arc, and in this segment, are all in the horizontal meridian.

Some of the high frequencies present in the video signal of Figure 10 were filtered out in the recording of Figure 12. This makes the fundus



TC-720522-19

FIGURE 10 VIDEO SIGNAL AND COMPUTED POSITION OUTPUTS FOR VOLUNTARY EYE MOVEMENTS



TB-720522-20

FIGURE 12 FILTERED VIDEO SIGNALS AND COMPUTED POSITION OUTPUTS FOR INVOLUNTARY EYE-MOVEMENTS

than these open-loop plots, but perhaps there are fundus motions which are not transmitted to the front of the eyeball.

We also printed out from the computer the entire video master frames used in Figures 10, 11, and 12, and calculated their power spectra. These calculations confirmed the effects of filtering, noise, and tracking accuracy discussed above. On the basis of these tests, we decided to proceed with the design and construction of the real-time, digital signal processor.

V DESIGN AND CONSTRUCTION OF SIGNAL PROCESSOR

The general outline of the electronic signal processor system is shown in the overall block diagram of Figure 13.

A. Scanning Rate and Sampling Rate

One of the basic design factors is the scanning rate, i.e., the number of scans per second. The higher the scan rate, the more often a correction is sent to the tracking servo. This argues for a high scan rate. However, a major limiting factor is the availability of mechanical drivers that can operate reliably at high speeds. A good compromise appears to be 400 scans per second, and for this purpose we use a drive motor rated at 24,000 rpm (i.e., 400 rps), though the motor can be operated at speeds as high as 30,000 rpm (500 rps).

Another major design parameter is the number of samples in each frame. Our simulation studies showed that 100 to 200 samples would be adequate from an accuracy point of view. Since we count photomultiplier pulses, however, as discussed below, other considerations favor a higher value, and we have chosen 512 samples for this design. The slower the sampling rate, the larger the number of counts per sample. The average light level with the present Hg-arc scanner yields approximately 1.6×10^7 photomultiplier pulses per second. At 400 frames per second, this means an average count of 40,000 per scan, and with 500 (actually 512) samples per frame, approximately 80 counts per sample (varying at most between 40 and 150 counts), which can be handled with 8-bit counters and registers. If the number of samples per frame were much less, we would need correspondingly larger registers to handle the larger number of pulses


counted per sample. This was the primary consideration in establishing the number of samples per frame.

Another important design consideration is the need for a relatively high-speed digital multiplier. It will be recalled that the basic algorithm forms and stores a master frame, and on each scan cycle cross-correlates the new signal with the stored master. This requires a single digital multiplication for each sample. With 400 scans per second, and with (approximately) 500 samples per scan, we have $400 \times 500 = 200,000$ samples per second, or a new sample every 5 microseconds. Assuming 8-bit binary numbers, this means that we require an 8-bit by 8-bit multiplication every 5 microseconds, which must include the time necessary to enter the numbers into the multiplier and transfer the resulting product to several different places. If we had chosen half as many samples per scan (i.e., 250), then we would have had twice as much time (8 microseconds) to accomplish the transfers and multiplication, but the multiplier would have had to be correspondingly larger in order to handle the larger binary numbers that result from a slower sampling rate (i.e., twice as many counts per sample). With presently available integrated electronic circuitry, it turns out that 5 microseconds is adequate time to accomplish this 8-bit by 8-bit multiplication. Thus the two main design parameters were chosen as 400 scans per second and 512 samples per scan.


B. Overall Timing

The overall timing of the machine is derived from a magnetic sync wheel mounted directly on the shaft of the scanner motor. This wheel has 256 teeth, from which we derive 256 sync pulses per revolution of the wheel. Intermediate sync pulses are generated by frequency-doubling circuitry, so that we obtain a total of 512 pulses per revolution (i.e.,



FOLDOUT FRAME



FOLDOUT, FRAME 2



FOLDOUT FRAME 3



**FIGURE 13 OVERALL SCHEMATIC DIAGRAM
OF SIGNAL PROCESSOR**

per scan cycle). These pulses are referred to as "minor sync" pulses. (It was only because of mechanical limitations that all 512 teeth were not machined directly onto the timing wheel.)

"Major sync" pulses are derived from a similar magnetic timing wheel, but occur only once per revolution. Thus, on each revolution, the signal processor receives one major sync pulse and 512 minor pulses. By deriving these sync pulses directly from the motor shaft, the subsequent signal processing is insensitive to variations in motor speed.

The source of these sync pulses is shown in the upper right portion of the overall block diagram, next to the block marked TIMING. However, various other timing signals must be generated as well. For example, as noted in Section I, we simplified the design by using a square-wave rather than a sinusoidal gating function in generating the X and Y displacement signals. For this purpose, the rotation cycle is broken up into four quadrants. The "sin" function is thus generated as a +1 function in the first and second quadrant of the rotation and a -1 in the third and fourth quadrants. The "cos" function is +1 in the second and third quadrants and -1 in the first and fourth quadrants. These gating functions are represented by the G_x , $\overline{G_x}$, and $\overline{G_y}$ signals.

C. Forming the Master Frame

The process of forming and storing the master frame can be seen along the top edge of Figure 13, starting with the photo-multiplier tube (PMT) at the left edge and ending with the CORE MEMORY at the right edge. To trace the formation of the master frame, let us begin with the basic process of forming the video signal itself.

For this purpose, we could start with an analog video signal, integrated directly at the PMT. But, for subsequent digital processing, this would then require an A-D (i.e., analog-to-digital) conversion

process. However, this would not be easy to achieve digitally at these speeds, and furthermore would represent loss of accuracy because of quantization noise. For these reasons, it was far more desirable to actually count individual pulses from the PMT and sample the count every minor sync cycle. This is the scheme in the present device. For this purpose we designed an 8-bit counter circuit that counts at a maximum rate of about 75 megacycles (which can be increased if necessary). At every minor sync cycle, the contents of this "photon counter" are read out and the counter is cleared.

The 8-bit number read out of the photon counter is entered into a register referred to as the 2δ -REGISTER. At every minor sync cycle, the contents of the 2δ -REGISTER are shifted one position and a new number is entered. This process goes on continuously, so that with a D-A (i.e., digital-to-analog) converter at any position along the register, we could create a (real-time) analog version of the video signal.

Recall that the master frame is actually a point-by-point difference between a retarded version and an advanced version of the original video signal. In effect, this forms a spatial derivative along the scanning path. It is this derivative signal that is stored as the master frame, and used for cross-correlation with each subsequent scan signal. The two time-shifted versions of the scan signal are generated from the 2δ register, which is 8 bits wide and 2δ bits long, where δ is the number of samples of advancement (or delay). The master-frame signal is obtained by differencing the 8-bit numbers stored at the first and last positions of the register. The result of this substitution is a 7-bit number, plus a sign bit. At the end of the first complete scan, we have obtained 512 samples comprising the master frame, which is stored for subsequent cross-correlation. The present system is designed so that δ can have one of four different values ($\delta = 5, 7, 9, \text{ or } 11$), which cover

the useful range determined by our simulation studies, as described in previous sections. Thus, if δ is set equal to 5, in effect a copy of the incoming sampled signal is made with a delay of $2\delta = 10$ samples, then subtracted point-by-point from the incoming signal and stored as the master frame.

As seen along the upper edge of Figure 13, individual pulses from the PMT are counted in the high-speed photon counter, and 8-bit numbers are sampled from the counter once per minor sync and entered into the 2δ -register. The first and last positions of the 2δ -register are subtracted from each other and the (7-bit-plus-sign) result is entered into the CORE MEMORY, which finally stores the 512 numbers that comprise a sampled version of the master frame. (The core memory used in the signal processor is manufactured by Fabritek Corp. and actually has a capacity of 1,024 8-bit words.)

Each time the START button is pushed, a new master difference frame is entered into the memory, through the process just described. Because the moment of pushing the button is unsynchronized with the basic machine timing, the circuitry is arranged so that the response to the button does not occur until the next major sync pulse. Then the write circuits of the core memory are activated during the following frame (major sync cycle).

D. Tracking Mode

Once the start button is released, the system automatically enters the tracking mode. In this mode, the core memory goes to read-only status. During each major sync cycle, the incoming video signal (derived from the center location of the 2δ -register) is multiplied point-by-point by the contents of the core storage (i.e., by the master difference frame). The desired multiplication is achieved by a standard

shift-and-add process. For this purpose, the 7-bit (sign excluded) number from the memory is entered into the right-shifting 7-bit register, shown at the center of the left edge of Figure 13. The current 8-bit binary number from the 28-register is entered at the right edge of a left-shifting 15-bit register. During the subsequent multiplication process (which must be completed during one period of the minor sync rate), the 15-bit register is shifted to the left one position at a time, while the 7-bit register is simultaneously shifted to the right. If during any shift, the last bit of the 7-bit register is a one, then the current contents of the 15-bit register are added into a 15-bit accumulator; if the last bit of the 7-bit register is a zero, the 15-bit register is simply shifted without being added. In this way, the final product is accumulated through a series of partial products. The partial products are accumulated by the joint action of the 15-bit adders and a 15-bit register. During the process, the current contents of the 15-bit accumulating register (which holds the current partial-product sum) is added to the incoming 15-bit number from the left-shifting 15-bit register via the adders, thus up-dating the partial-product sum.

In this way, one word of the current video signal and one word of the master frame are multiplied at each minor sync cycle. These individual products must then be accumulated over an entire major sync cycle, at the end of which we obtain a new reading for X, Y, and T, as discussed earlier in connection with the basic cross-correlation algorithm. The circuitry for properly accumulating the partial products for X, Y, and T is shown in the three rows at the lower right in Figure 13. Note that the output from the multiplier is always positive because the 8-bit number from the 28-register is always positive, and only the 7-bit number (sign excluded) from the memory is used. The sign bit from the core word is used, subsequently, in conjunction with the quadrant signals G_x , $\overline{G_x}$, G_y , and $\overline{G_y}$ described in Section B, to tell the X, Y, and T circuits whether to treat any particular product word as positive or negative. This control is accomplished by the 15-bit complementers at the inputs

to the X, Y, and T accumulator circuits. The actual accumulation is accomplished in each case by the joint action of a 20-bit adder and 20-bit register, the contents of the register being added to the new incoming product word (treated as a positive or negative number, depending on the core sign and the quadrant signals), and then reset into the 20-bit register. I.e., these 20-bit registers are up-dated on each minor sync cycle.

At each major sync-pulse, the contents of each 20-bit register are transferred into an 11-bit register which then holds this latest measure of X, Y, and T during the entire next major sync cycle, while a new value is being accumulated. We had calculated that an 11-bit number would represent sufficient accuracy for the desired servo drive characteristic. Initially, however, we did not know exactly how many bits of data we would actually obtain (this depends, among other things, on the basic photon rate). Therefore, we made provision to adjust the level within the 20-bit register from which these 11 bits are acquired for the servo output.

The 11-bit numbers that represent the current values of the X, Y, and T position signals are converted by three D-A units to analog signal representations, and these signals are then used for operating the servo mechanisms that control the position of the optical fiber in the Hg-arc scanner assembly.

E. Align Mode

In order to have as little interaction as possible between the X, Y, and T measurements, it is desirable to select regions of the retina in which there is approximately the same signal power in each quadrant of the scan circle. (To take an extreme example, if there were only one blood vessel, and it crossed only the top of the scan circle, we

could not distinguish between X and θ movements.) To help in aligning a new subject (i.e., finding a suitable location for the scanning circle on his retina), we have designed a special signal-processing mode that automatically provides a measure of the "signal content" in each quadrant of the scan. This mode takes certain measurements on the stored master frame and generates four meter readings, corresponding to the signal content in each quadrant. If the quadrants are not satisfactorily balanced, we can then move the scan circle to a new location and press the button for a new master frame.

To achieve this measure of signal power, we read out the master frame, which is stored in the magnetic core memory, and square it word by word, successively feeding each word directly into both inputs to the multiplier. This is achieved by the RUN/ALIGN switching box, which in the align mode feeds the core output, instead of the video signal, into the 15-bit left-shifting register. I.e., in the align mode the right-shifting 7-bit register and the left-shifting 15-bit register receive the same input, directly from the core memory. In this mode, the output of the multiplier is sent through a D-A converter and fed into an analog switch which segments the squared signal into four quadrants, for power measurement or direct viewing on a scope. In the RUN mode, the same D-A converter receives the video signal from the center position of the 26-register, so this may be viewed continuously during operation.

VI PRESENT STATUS

Essentially all our efforts during the past several months have been devoted to the construction, debugging, and testing of the signal processor. This device, which has about the same complexity as a small digital computer, now works in accord with the design shown in Figure 13, without making errors as far as we can determine. We connected it to the rest of the system as shown in Figure 1, and conducted open-loop tests with the model eye. At this point, we found that the design of the signal processor must be modified before we can obtain proper operation of the fundus tracker.

Analysis of our test results shows that there is a nonlinear interaction between the quadrant-gating signals (i.e., G_x , G_y , etc., in Figure 13) and the video signal, which results in a significant dc bias on the servo-control outputs. That is, even though the eye has not moved since the master-frame was formed, a large, steady X or Y output can occur. The expected eye-position signal is superimposed on this dc level, but of course we cannot close the loop in the presence of such a bias. The bias obtained varies with the location of the scanning circle on the fundus. But in the present design, the nature of the problem is such that wherever the best servo characteristic is obtained for X, the worst bias occurs in Y, and vice versa.

This bias occurs only when the output of the multiplier (i.e., the product of the video signal and the master frame) is correlated in some way with the quadrant-gating signal. Thus the square-wave gating signals in the present design are the main cause of the bias, because the horizontal and vertical "edges" of these gates generate the most bias when

they interact with horizontal and vertical edges in the video signal (which are most desirable for tracking).

This problem went undetected in all our computer simulation experiments because the computer program used sinusoidal gating functions as shown in Figure 3, not square waves. This program was written some years ago, before the signal processor had been designed (in fact, before we had decided on an all-digital scheme). Thus in spite of the tests described in Section II and III, we did not suspect that the exact shape of the gating functions would prove to be such a critical design feature.

It is also significant that the fundamental frequency of the video input (i.e., the frame rate) was suppressed in our computer simulations, in order to fit within the A/D amplitude range of the SDS-910. Thus the video input was completely uncorrelated with the ($\sin\theta$ or $\cos\theta$) gating signals. (Even in this case, the multiplier could have generated a $\sin\theta$ component by intermodulation; e.g., as the lower side-band of $\sin 2\theta$ and $\sin 3\theta$. But fortunately, the simulation experiments indicate that these cross-product terms will probably be negligible if we eliminate the fundamental bias.)

We can readily revert to sinusoidal gating functions, simply by modifying the control circuitry in the lower right part of Figure 13. The output of the digital multiplier must first be D/A converted and then analog-multiplied by sine and cosine functions (derived from the 400-cycle power supply of the scanning motor). Integration of these analog products would then provide the desired X and Y position outputs.

It would be more difficult to eliminate $\sin \theta$ components from the video signal. Because of our photon-counting input, this signal does not exist anywhere in analog form, and hence cannot be filtered directly. However, the analysis mentioned above shows that almost equally good

results can probably be obtained by eliminating the dc component of the video signal before multiplying, and this can be done on a digital basis. For this purpose, a subtractor is required between the output from the center position of the 2 δ -register and the input to the 7-bit right-shifting register of the multiplier. The second input to this new subtractor should be a binary number representing the dc level of the video signal, which can be obtained automatically by merely A/D converting the dc level of the video. Aside from certain higher-order terms, this will have the same effect as a frame-rate filter in removing any remaining correlation with the gating frequency.

Both of the design changes described above are relatively straightforward, but they must be implemented in order to achieve proper operation. Once they have been made, conversion of the video signal to position information should be essentially optimum.

It is not possible to make these design changes under the present contract, due to lack of funds. However, we intend to continue the development of the fundus tracker, and will seek other support if NASA Ames is unable to continue its sponsorship.

BIBLIOGRAPHY

1. Cornsweet, T. N.: Stabilized Image Techniques, in Recent Developments in Vision Research, NAS-NRC Vision Committee Publ. 1272, pp. 171-184 (1966).
2. Riggs, L. A., and A. M. L. Schick: Accuracy of retinal image stabilizing achieved with a plane mirror on a tightly-fitting contact lens, Vision Res. 8, 159-169 (1968).
3. Kelly, D. H., and H. D. Crane: Final Report on Phase I: Research Study of Fundus Tracker for Experiments in Stabilized Vision. SRI Project 6319. Contract NAS2-3995. November, 1967.
4. Kelly, D. H., et al.: Improved Retinal Image Stabilization Techniques. SRI Project 6319. Contract NAS2-3995. Qtr. Report 1, Phase II, 18 January to 31 March 1968.
5. *ibid.*: Qtr. Report 2, Phase II, 1 April to 1 July 1968.
6. *ibid.*: Qtr. Report 3, Phase II, 1 July to 1 October 1968.
7. *ibid.*: Qtr. Report 4, Phase II, 1 October 1968 to 1 January 1969.