

A HELMET MOUNTED SIGHT SYSTEM

Michael D. Stephenson and David Fosberry

SUMMARY

The role that a helmet mounted sight plays in modern fighter aircraft is described. Its capability in some situations has advantages over head up displays and these are described in some detail. The implementation of the systems using exposure control, analogue storage and signal detection are explained and how these functions are achieved using CCD technology to obtain a small, compact device for avionic use. Much of the latest work has been carried out using CCD devices manufactured at GEC Hirst Research Centre.

INTRODUCTION

Currently, pilots use a Head Up Display to accurately deliver weapons and stores to a desired target. The HUD consists of a semi-reflective surface which combines a carefully arranged CRT image with the outside world. This image is collimated at infinity so that when the pilot looks through the combiner glass he sees an image of the CRT symbology overlaid on the outside world.

The Head Up Display provides navigational and weapon information. For flying, he has the horizon bar and altitude scales etc. and for weapon delivery this can be combined with weapon trajectory vectors to achieve accurate weapon delivery. The Head Up Display (HUD) has excellent accuracy (of the order of 1mRad), but has the disadvantage that (i) when the pilot moves his head away from the HUD he loses vital information and (ii) to deliver a weapon he must hold the target within the field of view of his HUD (generally about 12°).

A new approach is to use a HUD which is mounted on the pilot's helmet, where the image is projected onto the visor. Using this method a pilot sees an image projected at infinity wherever he moves his head. (See Fig. 1.1). However, to be of any use in such roles as weapon delivery, the pilot's head position with respect to the airframe must be known. This has led to the development of the Helmet Position Sensor. Marconi Avionics examined various technologies and decided for a multiplicity of reasons, to choose an optical solution which led to the development of the Helmet Optical Position Sensor (HOPS).

Flight Automation Research Laboratory, Marconi Avionics Ltd.,
Airport Works, Rochester, Kent. MEL 2XX.

The helmet mounted display consists of a matrix addressable display of light emitting diodes (32 x 32). These are projected via a prism system onto a combiner glass in the pilot's visor. Here the image is combined with the outside world. The LED image is effectively focussed at infinity so that the symbology appears overlaid on the outside world (See Fig. 1.2). In modern designs the visor is shaped so that any part of the area that can pass in front of the pilots eye forms a combiner area.

The LED array is manufactured by a special process developed at GEC Hirst Research Laboratories. The emitting diodes are made with a brightness of 5000ft/Lamberts. The array is driven by hybrid circuits mounted on the helmet so that a minimum of wires are needed to supply the video data to the helmet.

The helmet mounted display is used to provide essential flight and weapon information to the pilot. It can also be used for weapon aiming or ground target designation and a reticle is provided for this purpose.

THE HELMET OPTICAL POSITION SENSORS

The purpose of the Helmet Position Sensor is to determine the direction of view of the helmet. The choice of the optical solution was decided by the then newly invented, Charge Coupled Device which, being solid state, offered a small size and optical accuracy suitable for cockpit mounting. The configuration chosen to implement this system (as shown in Fig. 2.1) consists of two isosceles triangles of high power infra-red LEDs (the LEDs are at the corners of the triangles) on each side of the helmet, viewed by two cockpit mounted 'V' slit cameras.

There are four possible combinations of LED triangle and camera, which enable accurate determination of roll, pitch and yaw, which fully specify helmet orientation over a large range of helmet positions and orientations. The LEDs in the triangle can be of any size, but the greater the LED spacing, the higher the accuracy of the system. The LED triangles can be of any size, but the greater the LED spacing, the higher the accuracy of the system.

The focal length is governed by the distance between the 'V' slit and the CCD. The images on the CCD are therefore a function of bit spacing of the CCD (which is defined extremely accurately) and the focal length, which can be set to anything desired. In practice, the acceptance angle of the camera is between 40° and 60°. This results in the position of the LEDs being measured accurately with respect to the camera axis in both X and Y directions.

The resultant image is gaussian shaped, this being set by a combination of the 'V' slit image shape and charge smearing in the CCD Bulk. However, the image is always well defined and does not change shape for even large deviations from the prime axis of the camera.

THE CAMERA

The camera uses a 1728 bit CCD as the heart of all the system operations. This device was chosen as it was the longest available device with a suitable architecture. A single device is used to detect movement in both x and y directions. To achieve this, an optical 'V' slit is used. The operation is shown in Fig. 2.2.

The 'V' slit projects two fan beams of light across the CCD. When the light source moves up and down the images move apart or together giving the 'y' direction movement. When the light source moves to the left or right, the two images move to the right or left but the spacing between the images remains the same.

EXPOSURE CONTROL

The cameras are generally fitted low in the cockpit so that the pilot's view is not restricted. At high altitude and in various parts of the globe, the sun intensity reaches a high level. This, combined with the position of the camera means that the CCD inevitably has to cope with full tropical sun. This requirement led to the investigation of an exposure control mechanism using standard CCDs. It is possible to use a neutral density filter in front of the CCD to reduce sunlight but this would also reduce the level being collected from the LED. A small amount of LED power is already being lost by a narrow band filter centred on the LED wavelength and any further reduction is not desirable.

The exposure control is used to increase the LED to Sun ratio. The LEDs used, can be driven at up to ten times their DC rating for short periods, with fairly low mark/space ratios. In this way, the camera can be exposed to a quite high LED illumination during a short exposure period, thus increasing the ratio of LED power received to sunlight power received. This improves the LED detectability and therefore optimises the range.

Also, the exposure control is used to ensure that the camera output is never saturated by the incident light, so that no information is lost, and clutter subtraction can be carried out to detect the LED images in any position.

The scale of the problem can be seen from the following figures. The intensity of solar radiation in free space at the wavelength of interest (the pass band of the filter) is about 0.8mW cm^{-2} per 100Å. A more detailed calculation (by multiplying the solar spectrum by the filter transmission

spectrum) gives a total of 5.667 mW cm^{-2} reaching the CCD. After correcting for the CCD responsivity spectrum, this gives an effective level of 3.69 mW cm^{-2} . Tests completed both in the laboratory and in flight trials at high altitude, indicate a saturation exposure of $0.3 \mu\text{J cm}^{-2}$. To avoid saturation, it can be seen that an exposure time of $81 \mu\text{s}$ or less is required.

Exposure control is achieved by sacrificing half of the output information. Normally, all odd photosites are read out via one CCD register, and the even photosites are read out via the other shift register. (See Fig. 2.4). In exposure control, one of the registers is used as a charge dump, or exposure drain, meaning that only half of the photosites can be output.

This is achieved by holding the exposure drain gate (ϕ_{xB}) high (open) at all times other than the exposure period. This means that all unwanted charge formed in the photosites is immediately drained away. When the exposure drain gate is closed, signal can then accumulate in the normal way, until it is transferred to the A register, after which ϕ_{xB} is raised again, see Fig. 2.5.

As can be seen from Fig. 2.4 the two registers feed into a common output amplifier. It is necessary to clock the exposure drain register on a two phase structure such as the CCD 121H, otherwise the barriers in the register stop charge reaching the output gate. See Fig. 2.6. It is also required that the charge from the two registers does not mix at the output.

The timing and drive of the waveforms can be arranged such that these conditions are satisfied and the charge from the exposure drain register is sunk, via the reset transistor to VDD, rather than being observed at the output.

This is achieved as follows (See Fig. 2.5). ϕ_{2B} is held at clock high potential at all times, and ϕ_{1B} is clocked such that it only goes high when ϕ_{2A} is low, the rising edge is not before the ϕ_R rising edge, and the falling edge is before the ϕ_R falling edge. This is achieved by raising ϕ_{1B} when the even reset ϕ_R is normally operated, and extending the ϕ_R period across the normal odd output period. This achieves the aims of blanking out the charge from the B register, and ensuring that it does not corrupt the signal charge in the A register.

It was found that, with the high illumination levels encountered, draining of the photosites during operation of the exposure control gate (ϕ_{xB}) was incomplete and also charge was found to be blooming along the device, through the B register. This was rectified by raising the B register clock high voltage (the voltage used for ϕ_{1B} , ϕ_{xB} and ϕ_{2B}) to 12V, and reducing the photogate voltage to 5V. These changes provide sufficient potential gradient to ensure efficient

removal of unwanted charge from the photosites, and provides larger potential wells to accommodate the charge in the B register. No harm is done to the performance of the photosites, although a certain amount of mixing of charge occurs into the selected bits (usually the odd bits), from each immediately preceding bit. This produces an enhanced sensitivity, and moves the centroid of each photosensitive area, but does not do significant damage to the resolution.

In other applications, if the image does not vary significantly from frame to frame, the loss of half the bits need not occur, since any register can read out either the odd or even bits, and can thus alternate between them, and they can then be interlaced off chip. However, in this application, the resolution is limited by charge crosstalk in the silicon, due to the great penetration depth of infra red light, and the signal can be reconstituted to the original 1728 bits or more, by interpolation techniques.

The integration period is between scans, and can therefore be varied over a very large range (10ms to $1\mu\text{s}$). The limitations at the long exposure end are due to dark current (an important area in the military temperature range) and at the short exposure end due to propagation times along the electrodes. A more realistic limitation to exposure time is that if the camera is exposed to light bright enough to necessitate $5\mu\text{s}$ exposure period, the exposure drain will not be able to handle the current coming from the photosites.

In our application, the exposure period is set at about $100\mu\text{s}$, which is sufficient to ensure that the solar illumination does not cause saturation. This differs from the calculated $81\mu\text{s}$ due to canopy absorption.

SUNLIGHT SUBTRACTION

To distinguish the LED signal from the sunlight, one must either remove the sunlight, or identify the signal in some other way. The pilot's head is in constant motion and therefore the LED signal cannot be considered periodic. This eliminates most of the normal correlation techniques and leaves us with rather brute force methods for detection.

Various methods of sunlight subtraction have been investigated. They all involve the following method. A scan of sunlight and reflection is taken and stored. This is then inverted and added to the next scan which contains sunlight, reflections and LEDs. This results in a combined video scan containing only the LEDs.

The first method of storage used Bucket Brigade delay lines, but unsatisfactory results were achieved, mainly due to transfer inefficiency. In the second attempt a fast A-D converter was used to digitise the CCD signal and subtraction was then carried out using logic. This was both expensive

and hardware intensive and resulted in a system which was commercially unsatisfactory.

Technology changes then pointed to the CCD as the state of the art analogue shift register and the function required to carry out subtraction is essentially that of a shift register. It was soon realised that with proper utilisation this could be achieved with the shift register on the imaging CCD.

The shift registers have fully accessible input ports, into which the signal from the 1st (background) scan is fed, after being passed through an amplifier with the appropriate gain, offset and delinearisation to give a throughput gain of -1. If one ensures that the shift register drive is incremented by the correct amount (i.e. it must not be overclocked) the fed-back, inverted bits arrive back where they started in the shift register, in the form of an offset of 80-90% of full charge plus the signal charge (inverted and therefore negative going). The resultant charge is in the range 0 - 90% full charge. This uses the following:-

LED signals + offset = (offset-background)+(LED signal + background), where the first parentheses contains the fed back signal and the second parentheses contains the second scan photosite information.

When implemented with exposure control, only the A register input port is utilised. The input technique used is 'Fill and Spill'. This behaves as a sampling input, removing the need for a sample and hold amplifier in the feed back.

Since the CCD utilises buried channel technology, the input port is very non-linear. This explains the use of a delinearisation stage (the amplifier has a square law term in the feedback over part of the range). The feedback offset is necessary to produce the charge offset in the register, which enables a device capable of carrying charge of one sign (negative) to store information of the opposite sign, by being superimposed on a negative offset.

The resultant signal can be separated from this offset by the use of a DC restore system which refers it to 0V. A schematic of the system is shown in 2.7. The residual signal after sunlight subtraction can be seen in Fig. 2.8.

SIGNAL DETECTION

Various methods of detection have been investigated, but it was finally realised that it would be possible to ultimately integrate a CCD transversal filter on to the video chip itself. Current investigations have shown that the Gaussian LED pulse shape changes very little with its position in front of the 'V' slit. It is therefore possible to define a transfer function for the filter. This enables the detect-

ion of small signals where the LED to camera distance becomes large. Present investigations use a discrete CCD transversal filter in the camera head. Work is still progressing along these lines to improve interpolation.

PROCESSING

The output of the camera is first analysed to determine the peaks of the two gaussian images of each LED. From the mean position and separation of the images, the direction of each LED, with respect to the camera, is calculated. From these, the cosines of the angles between each pair of LEDs is determined. These are then used in the following set of quadratic equations:-

$$x^2 + y^2 - 2c xy - 1 = 0$$

$$x^2 + z^2 - 2a yz - 1 = 0$$

$$z^2 + x^2 - 2b zx - 1 = 0$$

Where a, b and c are the cosines described above and x, y and z are the scaled distances, assuming the sides of the LED triangle are unity, to each LED from the camera. The equations are then solved iteratively, using a Newtons approximation method. There are four solutions, but usually only two unique solutions. The software then has to decide from the positional history, which solution is the correct one. The number of iterations to achieve a 'fit' is usually less than 3 unless the triangle is lost from the sight of the particular camera in use. The computer programme then reverts to an 'out of luck' routine which rapidly re-acquires the triangular fit.

RESULTS AND CONCLUSIONS

The functions in the helmet position sensor result in a small efficient system which utilises the CCD to its full, to give numbers related to camera angles. These numbers are fed to a 16-bit microprocessor to give roll, pitch and yaw angles at about 20 - 50 iterations/sec.

The computer selects the LED triangle and camera set combinations. When one LED set is lost from the view of a particular camera, a new selection is made, based on the previous result before the loss of information occurred. This enables 360° yaw coverage and a wide range of roll and pitch angles.

The results obtained produce accuracies of better than 0.5° RMS over an extremely wide range of angles. This is well beyond the aiming accuracy of the pilot himself. Other inaccuracies are introduced by the canopy and the very bumpy ride he gets at 250ft altitude and Mach 0.8, when in fact his neck muscles response limits tracking accuracy in severe turbulence. Some early results are shown in Fig. 2.9.

Successful operation has been achieved with the helmet mounted sight under realistic flight testing. Implementation of the sight and position sensor on to the helmet are still a subject for discussion as the addition of any weight, however small, is considered to be undesirable. This is emphasised by the fact that modern pilot's experience 9g in tight turns and 10z can quickly translate to 7½ lbs.

The helmet sight system enables the pilot to rapidly acquire and designate a target and deliver weapons on a first pass attack well outside the range of the Head Up Display. This requirement is essential in modern warfare if he is to achieve success in his missions.

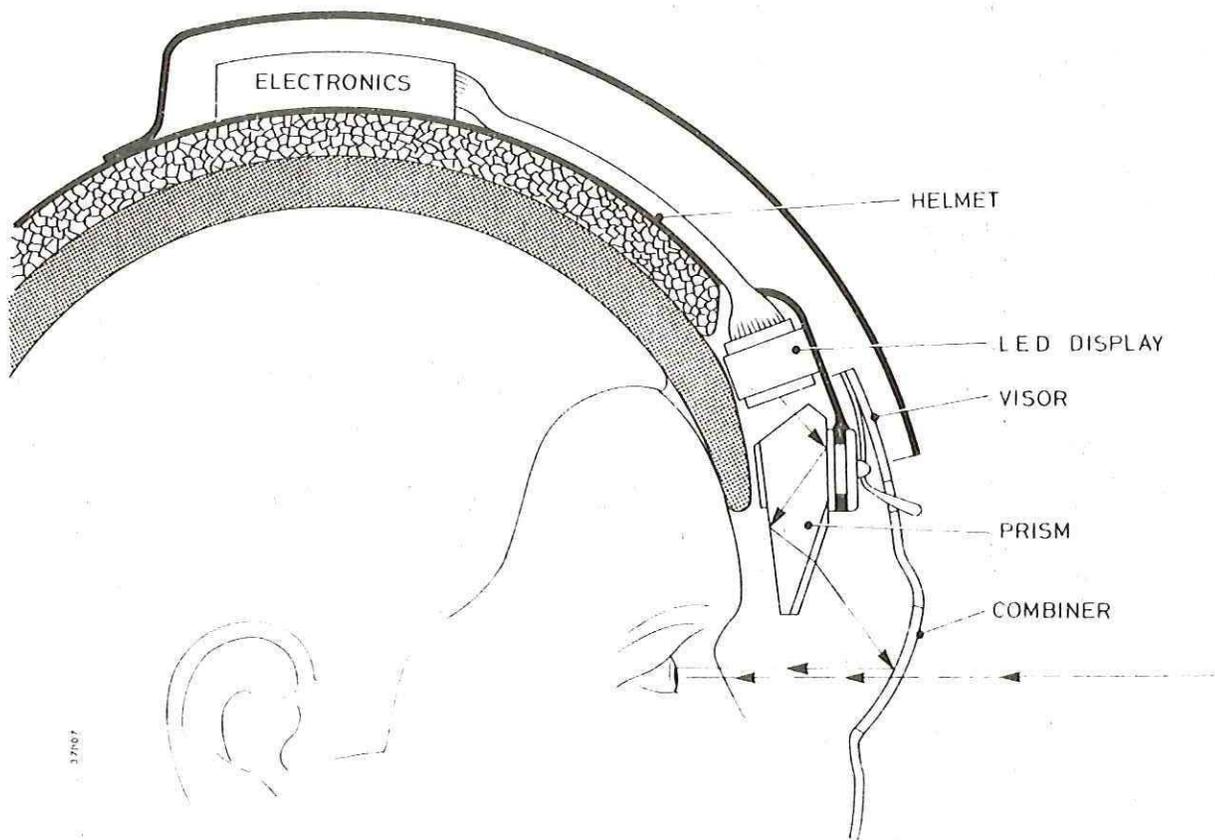


Fig. 1.1 Schematic Layout of Prototype HMD

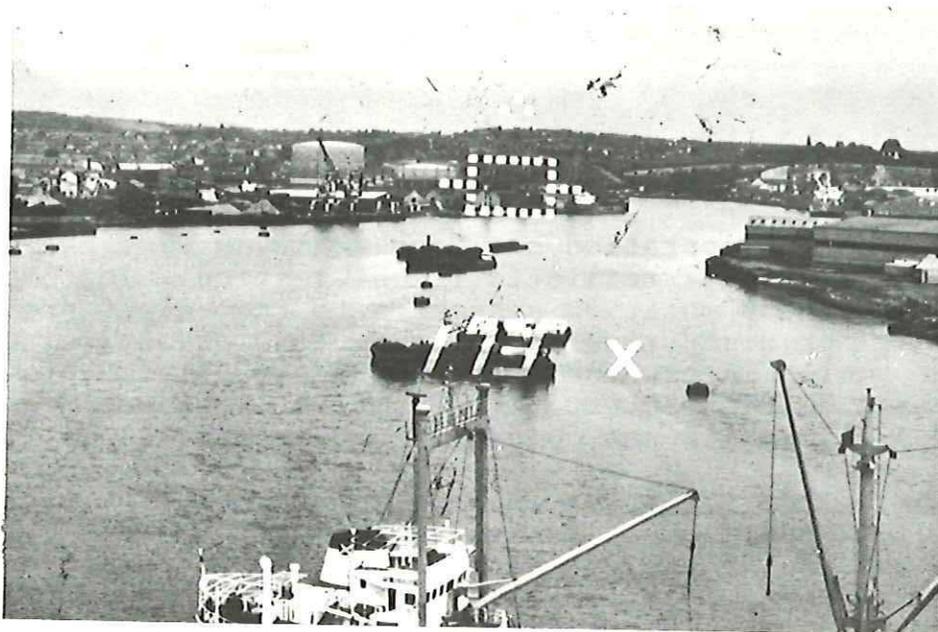


Fig. 1.2. View Through Helmet Sight

Helmet Optical Position Sensor System (HOPS)

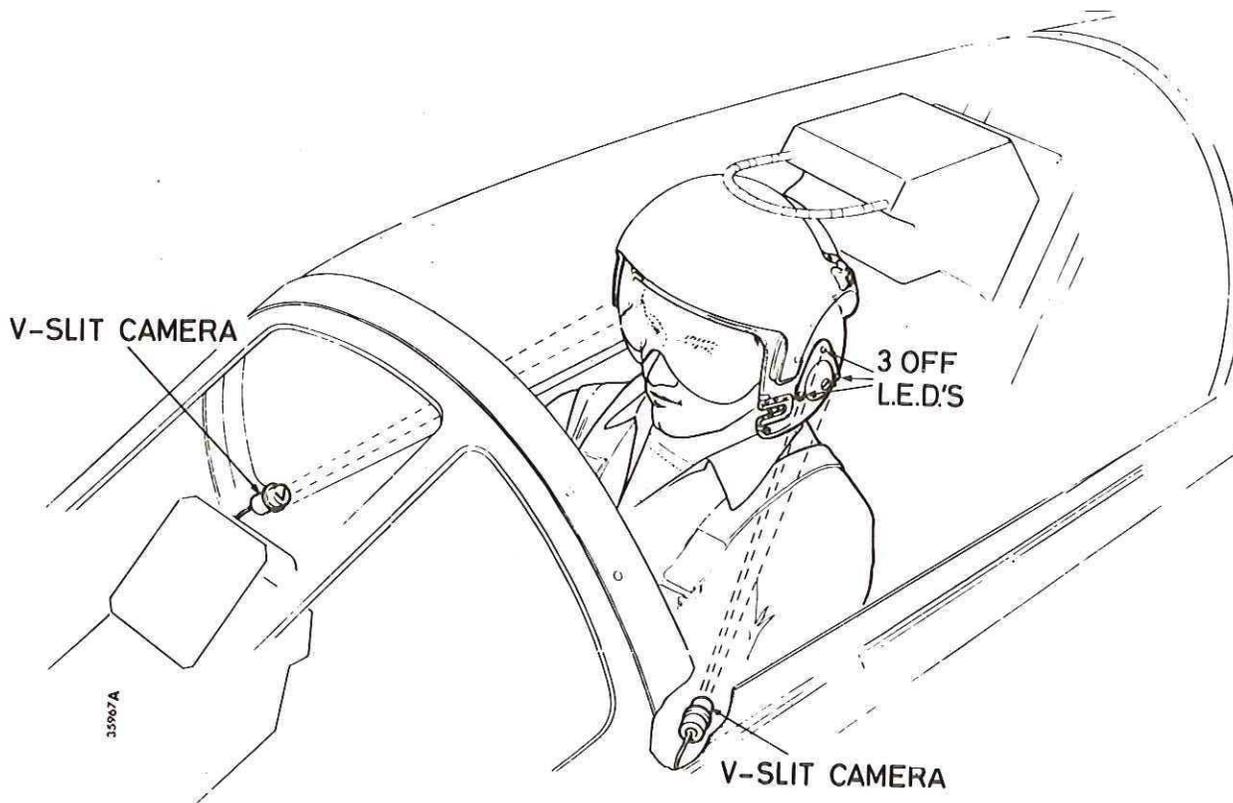


Fig. 2.1. Helmet Optical Position Sensor System (HOPS)

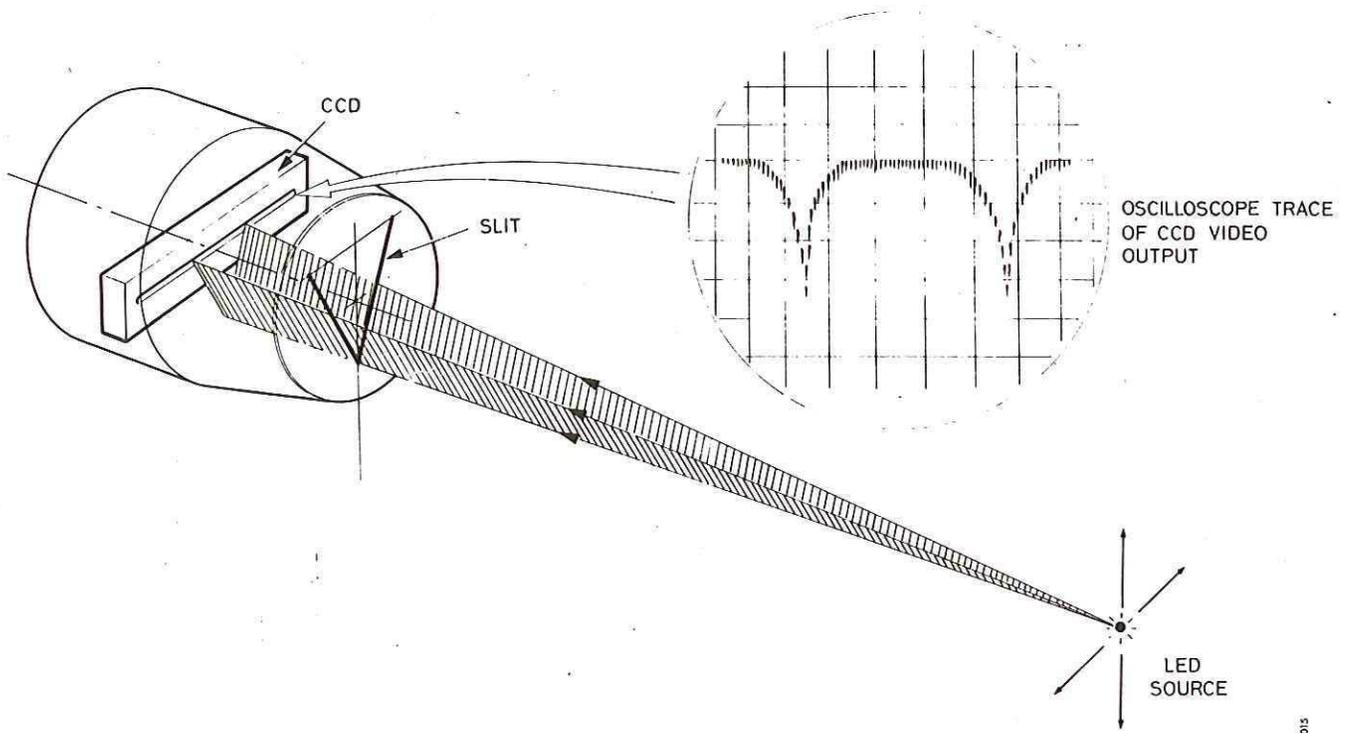


Fig. 2.2 V-slit Camera

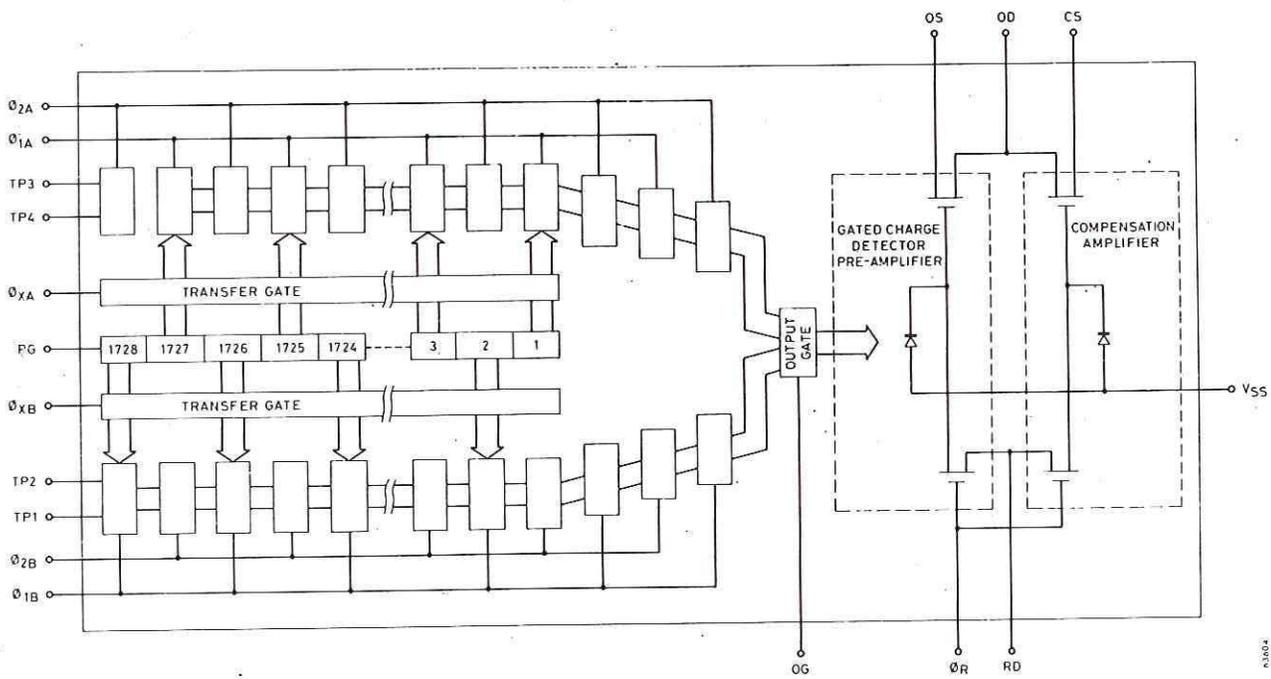


Figure 2.4. CCD Architecture

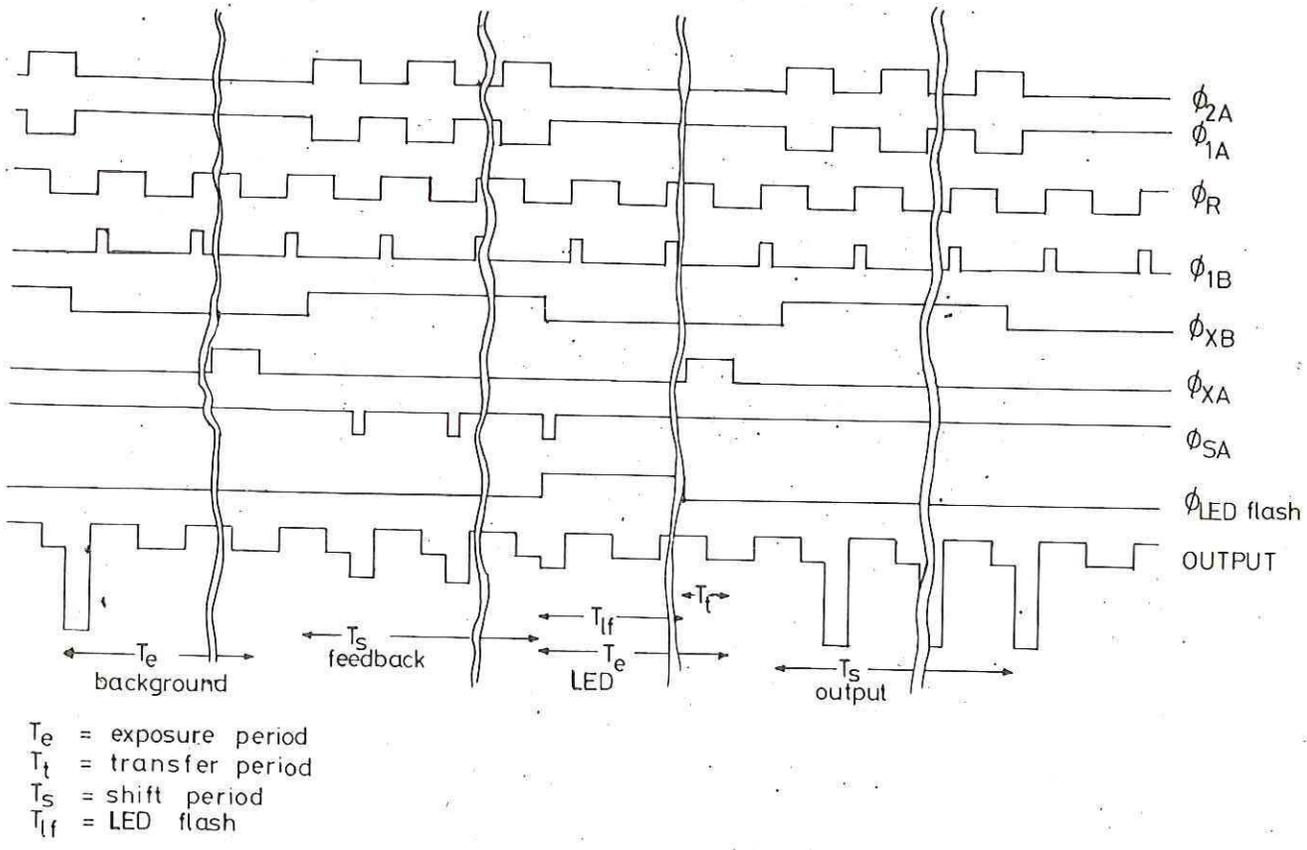


Fig. 2.5. CCD Drive Waveforms

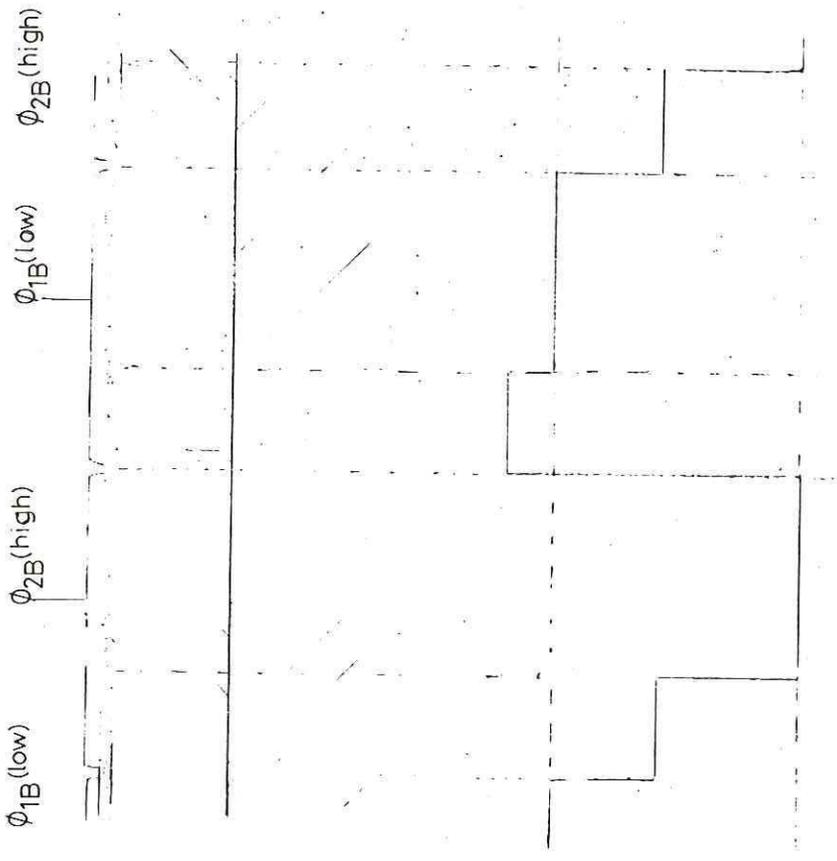


Fig. 2.6. 2 Phase Shift Register Potential Profile

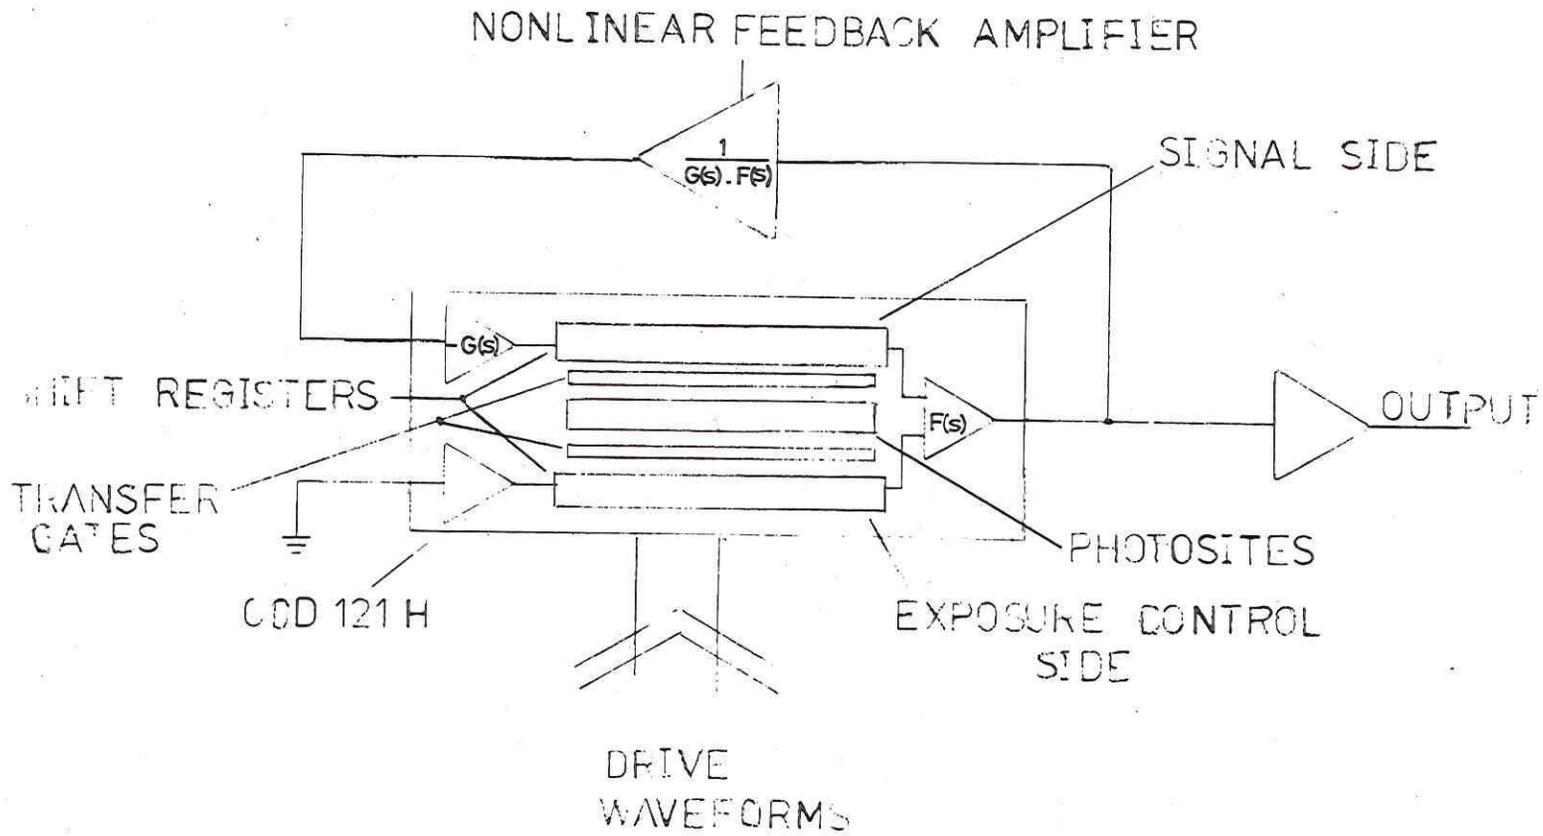
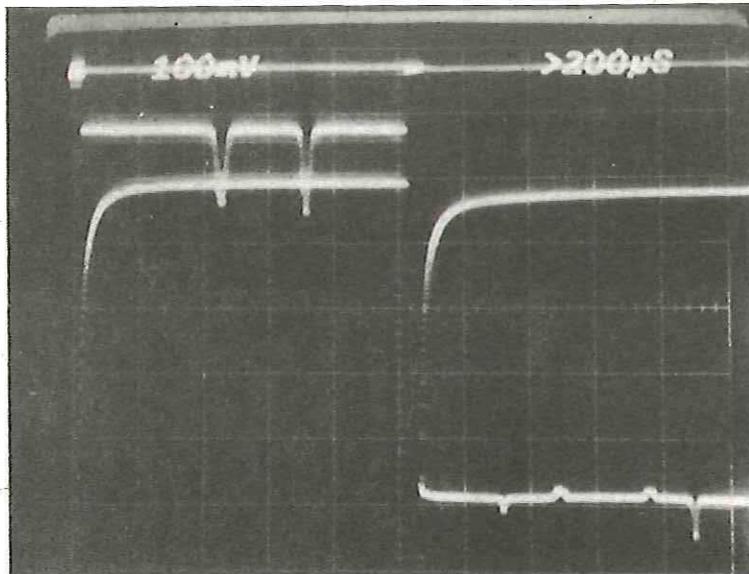


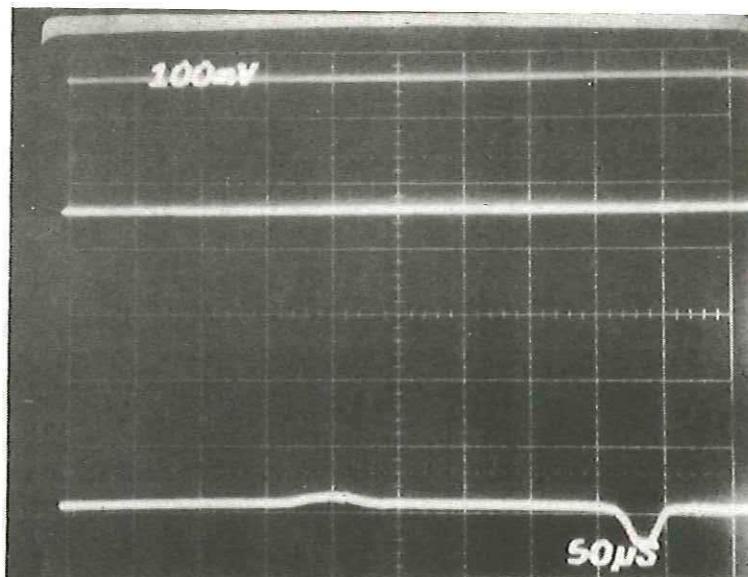
Fig. 2.7. Schematic of Subtraction Feedback

FIG. 2.8. Sunlight subtraction results

background and LED scans



expansion of LED scan showing
2nd sunlight residue and 2nd LED signal



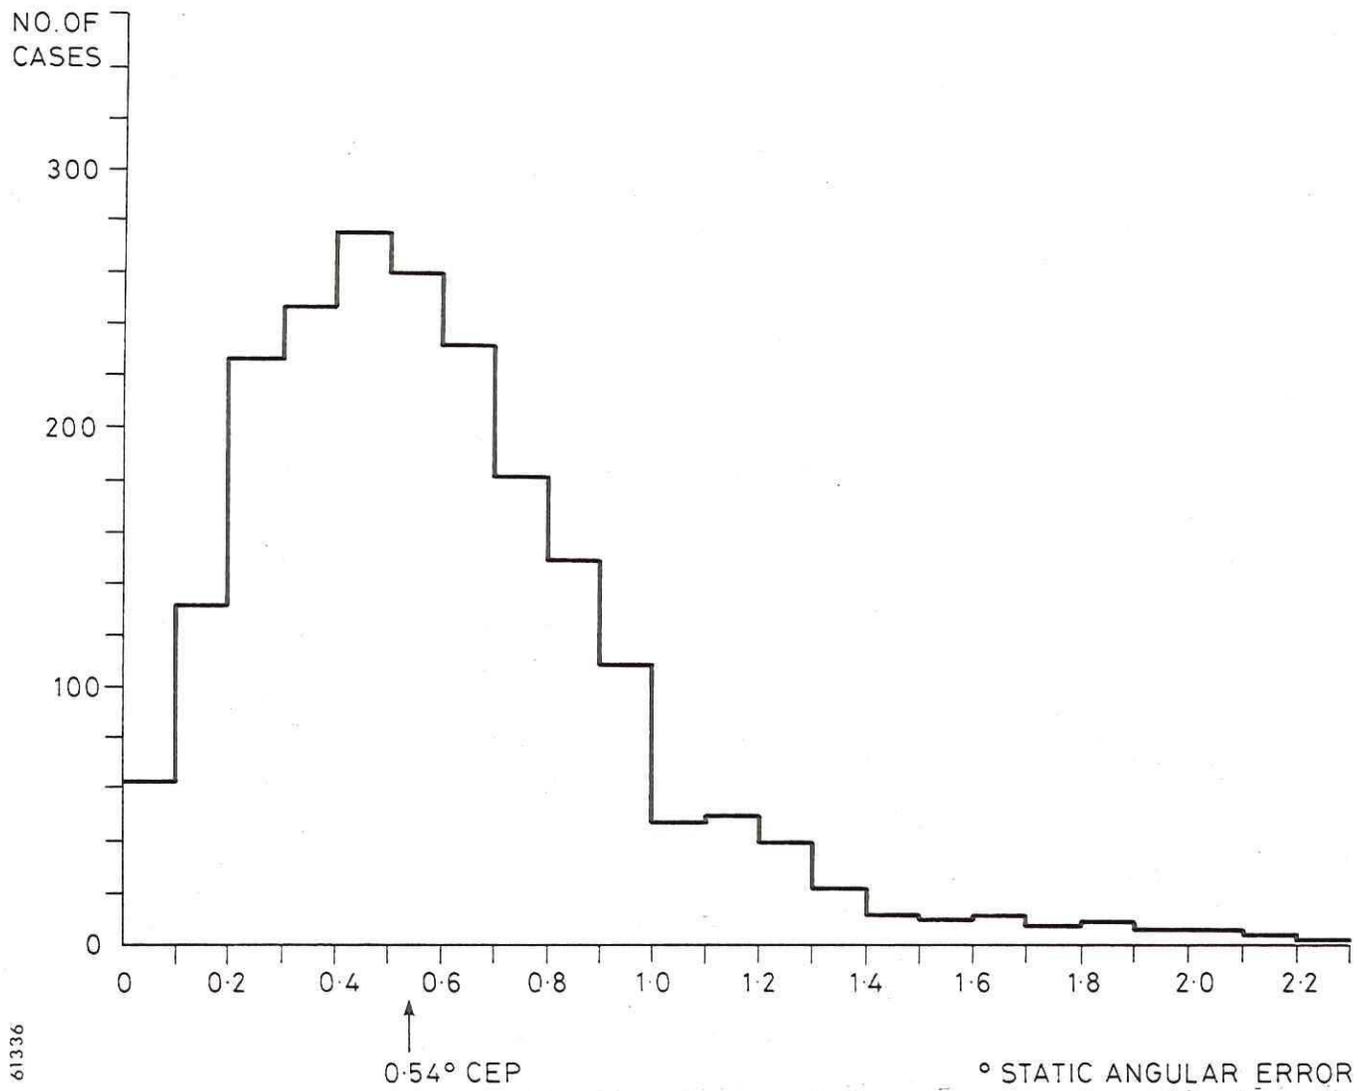


Fig. 2.9. HOPS Accuracy