UNIVERSITY OF CALIFORNIA

LICK OBSERVATORY TECHNICAL REPORTS

NO. 28

THE C.C.D DATA-TAKING SYSTEM

TOD LAUER
Tod Lauer

Santa Cruz, California

August, 1981

DUPLICATING CENTER REQUISITION
University of California – Santa Cruz

Complete all copies of form for each job. Attach white original, goldenrod and yellow copy to stencils or camera-ready copy and forward to the Duplicating Center in an envelope.

SEE INSTRUCTIONS ON BACK OF FORM

Date finished work required *(allow maximum time)*

FOR DUPLICATING CENTER ONLY	
Date Received	
Date Completed	
Negatives	
No. of Plates	Type
Total Impressions	

DESCRIPTION OF MATERIAL

Title and/or Form No. of Job: "Univ of Cal Lick Observatory Technical Reports No. 28, the CCD Data-Taking System"

[X] Letter/Form/Page	Finished Size 8½ x 11	Finished Quantity Required 30
[ ] Poster	Finished Size	Finished Quantity Required
[ ] Card/Envelope	Finished Size	Finished Quantity Required

OFFSET

No. Pages: 44 + cover	No. Copies Each Page: 30 + 5 extra front covers
Ink Color(s): black	Paper Color(s): white
Paper Type: stock	
[ ] Metal Plate	[X] Paper Plate
[X] Print One Side	[ ] Print Two Sides
Plate Retention [ ] Yes	[X] No

MIMEOGRAPH

No. Stencils (attach)
No. Copies Each
Paper Color(s)
Paper Type
Stencil Retention [ ] Yes [ ] No

CHECK REQUIRED ITEMS BELOW

STAPLE	✓	PUNCH	✓	FOLD	✓	PAD	✓
Upper Left Corner	X	Two (2) Holes		Single		Amount Per Pad	
Marginal		Three (3) Holes		Letter			
Two (2) Staples		Other		Other		Regular Glue	
Three (3) Staples						NCR Fan Apart	

SPECIAL INSTRUCTIONS

Please supply front and back covers in green stock like the sample attached.
The front cover will be printed, the back cover will be blank.

PLEASE PRINT UP 5 EXTRA FRONT COVERS.

Staple upper left corner with the staple placed parallel to the left edge.

Distribution By Mail To:	Send Extra Copies To:

Deliver To: Marilyn Oreck	Room 176	Bldg. Nat. Sci. 2	Pick-up (Call Ext. when ready)

Date: August 1981	Office Name: Lick Observatory
Requested By: Tod Lauer	Approved for Payment Authorized Signature: Marlene Couture

Name of Account To Be Charged: NASA NSG 7503 #S3

CHARGE ACCT CODING	LOC	ACCT	FUND	SUB	OBJ
	7	444645	23209	7	

Reference No.

$ TOTAL COST

UNIVERSITY OF CALIFORNIA

LICK OBSERVATORY TECHNICAL REPORTS

NO. 28

THE C.C.D. DATA-TAKING SYSTEM

Tod Lauer

Santa Cruz, California

August, 1981

TABLE OF CONTENTS

INTRODUCTION

This report describes the use and structure of the C.C.D. data taking system on the Lick Observatory PDP-8 computers. In particular, this report provides details on the procedures used to take data with the C.C.D., the use of the analysis programs, how the system operates, how the data are stored, and how the output devices are used. Reading the first section of the report, the "quick overview", will give you enough information to get started and operate the system under most situations. The other sections will give you a deeper understanding of how the system works, so you can use it with some amount of confidence and expertise. It is especially important that you read the section on the use of the 9 track tape drive, as this is the primary image storage device and can cause loss of time at the telescope if misused. Lastly, this report does not discuss the C.C.D.'s characteristics, the use of its controller, or how to reduce the image data in any detail.

The system is simple and straight forward to use. Its primary function is to control, read out, and save C.C.D. observations, which it does in a semi-automatic fashion to minimize user intervention. The observer merely configures the system as desired, and pushes the start button when ready to take data. The system will then preserve its configuration until changed by the observer. The system also has a small library of programs that an observer can use to display, manipulate, and examine the image data. These programs can be used at the telescope while taking data, so that the observer can evaluate the observations as they are made.

The system also contains several error detecting and fix-up routines to protect the observer from loosing data and time because of mistakes or confusion in using the system, or the unlikely occurrence of crashes in the system itself. While it's not possible to foresee all types of mistakes, modes of failure, or to protect the system against direct assault, a serious effort has been made to foresee and provide a solution to all problems that might be encountered while taking data or working with the programs. The hope is that the observer can use the system without fear, and can feel free to experiment with it as desired.

As it now stands, the system is designed to take direct images with a T.I. 500 by 500 C.C.D. As new chips become available, and different modes of observation desired, the system will be modified to handle them. Further, more programs will be added to the system as needed.

Finally, credit is due to the many people who provided useful suggestions and criticisms on the structure and use of the data taking system. Specifically, I would like to thank Lloyd Robinson for his ideas on the overall system philosophy and structure, Bob Kibrick (who remembered the line from "2001") for many informative discussions on the use of the tape drives, the area scanner memory, and too many other things to enumerate, and to Richard Stover and Diane Wooden who were the first ones to use the system and provided much help in debugging and smoothing out its operation.

A QUICK OVERVIEW OF DATA TAKING WITH THE C.C.D.

I. Procedure Introduction

This section gives a quick overview on how to set up, and use the C.C.D. data taking system. This outline is not intended to be a complete guide on how to use the system, but should enable you to get started and operate the C.C.D. under normal circumstances. For a more complete understanding of the system, you are referred to the other sections of this report.

This procedure assumes the normal configuration of a PDP-8 computer, a 9 track tape drive, and the area scanner memory. As of this writing, you must provide a TV monitor for the display of images, which should be patched with a 50 or 75 ohm BNC coaxial cable into the "TV" terminal on the back of the area scanner memory. This procedure also assumes that the C.C.D. and its controller have already been properly linked to the system. If the tape drive is not available, then data can be stored on floppy disks. If the area scanner memory is not available, then crude image displays can be generated on the C.R.T.

II. Before Bootstrapping

1. Load a magnetic tape onto the tape drive, making sure that the tape has a write ring, and that the drive is switched on line when the load sequence is complete.

2. Turn on the area scanner memory. Push the stop and reset buttons, and check to see that the unit is write enabled.

3. Turn the C.R.T. on, and allow it to warm up.

III. Bootstrapping

1. Bootstrap FOCAL and type "G", followed by a carriage return to load the system. (All FOCAL commands must be followed by a carriage return.) The system will check the C.C.D. and exposure status.

2. If an exposure is not in progress, or in need of readout, or a crash while writing to tape has not occurred previously, the system will respond with "INITIALIZE C.C.D.? <Y/N>". If no, the previous parameters will be preserved and displayed, and the system will return to the monitor mode. If yes, an initialization dialog will be started. If the C.C.D. controller is not properly linked up to the system, however, the message "PERFORM DATA REDUCTION ONLY - C.C.D. NOT AVAILABLE" will be printed out, and the system will return to the monitor mode.

3. If the initialization dialog is followed, the next step is to enter the U.T. time of day from the C.C.D. controller. This time will be used as a backup if other time standards are not available.

4. The dialog next prints the initial size of the C.C.D. and asks, "SET WINDOW? <Y/N>". If yes, the window size is read from switch 1,1 and the origin is entered from the keyboard. If no, the initial size is used. Note: The observing window parameters can be changed at any time later.

5. The dialog next asks, "INITIALIZE IBM TAPE? <Y/N>". (If it doesn't, then the tape is unavailable.) If yes, the tape will rewind, and you will be asked, "WRITE AT END OF OLD TAPE (O), OR START NEW TAPE (N)." If "O", the tape will simply advance to the end of data. If "N", you will be asked to confirm the new tape; on BRIGHT drives with <Y/N>, on CIPHER drives by specifying high (1600 BPI) or low (800 BPI) density. Specify high density, unless you have specific reasons for wanting low density. Note: This initialization procedure can be called at any time later, but must be called whenever new tapes are loaded.

6. Finally, the bootstrap dialog asks for an initial observation number, which must be between 0 and 4095 inclusive. This is part of the log information, and can be changed at any later time.

7. The dialog concludes by sounding a bell and lighting the ready light. The system is now available for commands.

IV. Before Starting an Exposure

1. To provide a backup output device for the tape drive, place a blank disk in unit (7), making sure the unit is write enabled. If the disk needs clearing, call the edit program, making sure switch 4,11 is raised to indicate the floppy disk to the editor.

2. If convenient, it is a good idea to call the scrub program to clear out any residual charge in the C.C.D. Set the scrub cycle time with the exposure time switches. Lock the calling switch in the up position for repeated scrubbing.

3. Call in programs from the switch panel as needed. All programs will return to monitor mode at completion. You can also escape any program, except those writing to tape, by typing a Control-C, followed by a "G" then carriage return, to return to the monitor.

V. Starting an Exposure

1. Set the exposure time with switches 2,1 and 1,7.

2. For a dark (shutter closed) exposure, raise switch 3,3.

3. If needed, enter the observation header by calling the "change header" program. This header will be preserved until you change it.

4. Set the output device. Although this is not necessary now, if the tape drive is selected, its status will be checked, and a warning message

printed if it is not ready to record data.

5. Push the start button (switch 3,8).

VI. During an Exposure

1. During an exposure, the system reads the exposure time counter and waits for commands. Programs can be called in, but the clock display will not be revised during their execution. Control will be returned to the exposure at their completion, or if "G" (then carriage return) is entered from FOCAL following a crash or Control-C. The C.C.D. controller does not generate interrupts, so it is your responsibility to make sure the system is in the monitor state at the expected time of readout. If the system is not in control then, the readout will be delayed until it is. The shutter will still close on time, however. Note: To remind you when the monitor is not in control of an exposure, a lockout light is lit during this time.

2. The exposure status and information are safe over bootstrapping.

3. To stop the exposure and force a readout, push the stop button (switch 3,10).

4. To abort the exposure and discard it, push the reset button (switch 3,9). The system will leave the monitor state, but no further action will be taken until you confirm your decision from the keyboard.

5. To modify the exposure time, depress switch 3,4. Enter and confirm the new exposure time from the keyboard. If the new exposure time is less than the elapsed time, the exposure will end, and the readout will be started.

6. To pause during an exposure, raise switch 3,6. This stops the exposure clock, and closes the shutter until the switch is lowered. During this time the pause light is lit, and programs can still be called in and out.

7. Exposure parameters such as C.C.D. window size, observation number, and the header, can be changed at any time during the exposure, however, it is best to make all changes beforehand if the exposure is short.

VII. Readout

1. The output device switch (switch 2,11) is read at the start of the readout cycle.

2. If the tape drive is selected, the first step is to hunt for the end of data. If the tape is positioned more than a few images away from the end of data, this hunt could take a significant amount of time. In this case, it is best to manually force the tape to the end of data, while the exposure is still in progress, by requesting the display of an image number larger than exists.

The system checks for tape errors as follows. If the tape drive has not been initialized yet, the initialization procedure is invoked at this time. If the tape is off line, a warning is sounded until it is switched on line. If the tape is not ready for any other reason, the system switches the output over to the backup disk. Note: Control-C is disabled during the readout to tape.

3. If the floppy disk (Unit 7) is selected, a check of the image size is made against the remaining disk capacity. If the image is too big to fit on one disk, (400 by 400 pixels is about the limit) a warning is printed at the start of the readout, and then when the disk must be changed to continue the data storage.

4. If the TV is selected, readout is to the area scanner memory. If the data taking window is larger in either dimension than 256 pixels, the image is compressed in both dimensions until it fits into memory.

5. If the "quick look" is selected, the full array is read out to the area scanner memory, with compression as needed. Note: This procedure does not affect the previous data taking window.

6. The readout proceeds.

7. If a write to tape is unsuccessful, or the end of tape is reached, output is switched over to floppy disk, and the last file on tape is closed. To prevent further errors, you should now unload the tape.

8. Errors on floppy disk will interrupt the readout until you replace the disk, or take other action such as checking the loading of the disk or the write status of the disk unit.

9. If a system crash occurs during the readout, restart the system by typing "G" (then carriage return). If the output has been to tape, the output file will be closed and you will be asked if the image readout should continue on floppy disk. On other devices, however, you will simply be asked if the readout should be continued.

10. At the end of the readout, exposure status information is printed out for all devices, except the "quick look". This printout can be bypassed by raising the mute switch (switch 4,8) beforehand.

11. If the output has been to 9 track tape or floppy disk, the TV memory will be erased, and the image displayed. If the no display switch (switch 3,6) has been raised beforehand, however, the display will be bypassed. If the output spans more than one floppy, the display will also be bypassed.

12. If the output has been to 9 track tape or floppy disk, the observation number will be incremented by one.

13. The system returns to the monitor state, ready to begin the next exposure or to bring in other programs.

DESCRIPTION OF SYSTEM PROGRAMS

I. Introduction

Below is a listing and brief description of the programs available with the C.C.D. data taking system as of this writing. This list should be regarded as provisional; existing programs may be modified and new ones added to the system as future need dictates. Nevertheless, there are already a wide variety of programs to assist you in taking data and evaluating your observations. Basically, the programs can be divided into four rough categories, those to set up the system for data taking, those to control the input and output of images, those to display the images, and finally, those to help you analyze the observations.

Programs can be called from the switch panel whenever the system is in its monitor state, whether an exposure is in progress or not. Simply set all controlling switches beforehand and call the program by depressing the appropriate calling switch. The system will leave the monitor state, and the program will execute. Control will return to the monitor at the program's completion. All programs, except those writing to 9 track tape, can also be escaped by typing Control-C, and then "G" (carriage return) to return to the monitor. Lastly, it must be remembered that the system will not be able to control or monitor exposures while a program is executing. It is your responsibility to make sure the system is in the monitor state at the end of all exposures so the readout procedure can be started.

II. Observation Set-Up Programs

Scrub:

This program is used to repeatedly readout or "scrub" the C.C.D., in an attempt to remove any charge retained in the chip's wells after an excessively bright or saturated exposure. The scrub program should be run after the C.C.D. has been mounted, when it has been exposed to dome lights, and between exposures when convenient. To scrub the C.C.D., set a cycle time with the exposure time switches, and call the program. The C.C.D. will then be read out after the cycle time selected. The choice of this time is not critical; it merely provides for varying time scales of retained charge leakage. For repeated scrubbing, lock the calling switch in the up position.

Focus:

This program helps you set up the C.C.D. for direct imaging by displaying successive images of some object in the chip's field, allowing you to test the focus of the telescope. To focus the telescope on the C.C.D., first find a field which contains a bright star (5 to 8 m) and read out the chip to the TV. (The "quick look" readout is recommended for its full view) Call the focus program, and select the star with the cursor. The screen will then be erased, ready to hold up to nine successive 64 by 64 pixel images centered on the part of the chip marked by the cursor. Set an exposure time on the switch panel and push the start button to read out the first focus test image. Examine the image, adjust the telescope focus as desired, and push start when ready to test the focus again. Continue as needed. The TV screen will hold up to nine images, allowing you to evaluate your progress. When the screen is full, the older images will be overwritten in the order in which they were taken. To escape the program, quickly tap (so as not to call other programs) the lower program calling switches. Note: It may be necessary to guide the telescope between exposures to keep the test star in the field of view. Further, it may be necessary to reduce the exposure time as the star's focus improves.

Set Telescope:

This is the standard Lick Observatory precession routine. It reads the time from the telescope controller or other time standards.

Set C.C.D.:

This program allows you to select which part of the C.C.D. chip is to be used for observations. Select a window size, or request one to be entered from the keyboard, with switch 1,1 and call the program. (The preset windows are square, but this is not a restriction.) If the window is to be entered from the keyboard, you will be asked for the number of rows and columns in the window. What ever the case, the program always checks the selected window size against the size of the chip, and its number of columns for even parity. If less than the full chip is used, you are then asked for the window's starting row and column number. (The chip's rows and columns are numbered from "0" to the size of the chip in the appropriate dimension less one.) Both numbers are then checked for consistency with the window and chip size; the starting column is further checked for even parity. If an error is detected in any entry, the programs will ask for that entry again.

Set Observation Number:

Calling this program prints the old observation number and asks for a new one to be entered from the keyboard. Its range of values must be between "0" and "4095". This number will be incremented by one every time an image is read out of the C.C.D. or TV into 9 track tape or floppy disk, and is part of every image's parameter list.

Change Header:

Calling this program prints the old observation header or label, and requests a new one from the keyboard. This header will be written to the image file of all observations read out of the C.C.D., and can contain up to 62 alphanumeric characters.

III. Data Input and Output Handling Programs

Save TV:

This program will write the image held in the TV or area scanner memory to 9 track tape or floppy disk. The default output device is tape; to transfer the image to floppy disk, the "data output" switch (switch 2,11) must be set to the floppy position. The program also generates a histogram record of the image pixels, and increments the observation number by one at its completion. Note: Raising the "mute" switch (switch 4,8) will suppress the image status printout at the completion of the program.

Initialize Tape:

This program is used to set up and position a new or old 9 track tape for initial or additional data storage. For details, see the section of this report on the use of 9 track tape. Tapes loaded for reading only do not need to be initialized.

Edit Output:

This program is used to delete image files stored on 9 track tape or floppy disks. To choose between the two devices, set the "data input" switch (switch 4,11) before calling the program. Because tape is a sequential storage device, only the last image on tape can be deleted without affecting other images; a particular file on tape can be deleted only by first deleting all files that come after it. To edit the tape then, the program locates the last image file stored, prints its header and parameters, and asks you to confirm the deletion. If you concur, the image is deleted. For further editing, the program can be called again. For floppy disks, the program will ask you if the whole disk is to be cleared or just the last image. As above, you will be asked to confirm any deletions before the program will take action.

Directory:

This program will produce a listing of the images stored on a reel of tape or floppy disk. To choose between the two devices, set the "data input" switch (switch 4,11) before calling the program. The directory program has two modes of operation controlled by the "mute" switch (switch 4,11). In the muted mode, the program simply prints a table listing each image's basic parameters and the first 25 characters of its header. In the normal mode, the program will ask you for an overall directory header and produce a paged output giving a much more extensive listing for each image stored.

Floppy to Tape:

This program will transfer images saved on floppy disk to 9 track tape. After the program is called, you will be asked if the whole disk is to be transferred (one image at a time) or just one particular image. If necessary, a histogram record is generated for each image during the data transfer.

Tape to Floppy:

This program will transfer individual images stored on 9 track tape to any floppy or dectape unit. After calling the program, you will be asked which unit the image is to be transferred to, and then which image on tape to transfer. If necessary, a histogram record is generated for each image during the data transfer. By raising the "window" switch (switch 4,9) before calling the program, only the section of the image which falls within the preset reduction window will be transferred.

Reduction Window:

This program loads a set of window parameters into core, which are then used by other programs to select a subsection of a stored image for display, transfer, and other operations. The reduction window is set up in an identical fashion to the observation window as described in the "Set C.C.D." program write up. The window is invoked by raising the "window" switch (switch 4,9) before calling any program which makes use of it.

IV. Image Display Programs

Display:

This program will display images on the TV screen by transferring data stored on 9 track tape or floppy disk to the area scanner memory. Choose between the two storage devices by setting the "data input" switch (switch 4,11) before calling the program. To display a preselected subsection of a stored image, raise the "window" switch (switch 4,9) before calling the program. (See above for setting up a window.) Lastly, to suppress a printout of the image's parameters, raise the "mute" switch (switch 4,8) before calling the program. After calling the program, you will simply be asked the number of the image you wish to display. Note: Remember that images larger than 256 pixels in either dimension will be compressed in both dimensions in order to fit into the display memory.

Process:

This is actually a set of four subroutines which can be used to enhance or highlight features in the TV display in various ways. Each of the subroutines works directly on the data held in the area scanner memory. Set the "process" switch (switch 2,3) to one of the subroutines listed below before calling the main process program.

Linear Stretch:

This subroutine can be used to boost the contrast of a displayed image by performing a linear "stretch" of the pixel intensity values. After calling the subroutine, you will be asked for a "low level" and "upper level" of pixel intensity values. Pixels with intensities within this range will have their intensities mapped in a linear fashion to the full intensity range of "0" to "4095". Pixels with intensities below the low level, and above the upper level, will be mapped to intensities of "0" and "4095" respectively.

Log Stretch:

This subroutine is designed to compress the dynamic range of the pixel intensities by displaying the logarithm of the intensities renormalized to the full "0" to "4095" intensity range. The main effect of this process is to greatly enhance the contrast between the low intensity features of an image at the expense of contrast between the brighter features. Indeed, a linear stretch is often needed after a log stretch to restore some contrast between the bright features of the image and its background.

Invert:

This subroutine will invert the display's gray scale, converting a "positive" image into its "negative" and visa-versa.

Isophote:

This subroutine can be used to break the display down into a series of regularly or irregularly spaced isophotes, as desired. After calling the subroutine, you will be asked to choose between the two options. If you select regular isophotes, you will be asked for an initial offset and the intensity spacing of the isophotes. Pixels with intensities below the offset level will be zeroed. If you select irregular isophotes, you will then be asked to specify the spacing of individual isophotes as desired. Pixels with intensities above or below all the specified isophotes will be zeroed. Note: It is possible to use this subroutine to zero out high valued noisy or extraneous pixels by breaking the image into one wide ranged "irregular" isophote containing only acceptable pixel intensity values.

C.R.T. Display:

This program can be used to generate a crude image display on the C.R.T.'s screen by displaying each pixel as a dot matrix. There are ten such matrices, hence ten gray levels available for discriminating between pixels of differing intensities. Select between storage devices by setting the "data input" switch (switch 4,11) before calling the program. After calling the program, you will be asked for a zero offset and the intensity interval between gray levels. Because of the small number of gray levels, it is probably best to specify a small interval between gray levels and let the image be broken down into isophotes, rather than trying to match the full dynamic range of the image. Note: If the image is larger than 200 pixels in either dimension (not 256 pixels as for TV images), the image will be compressed in both dimensions as needed in order to fit on the screen.

V. Data Analysis and Reduction Programs

Pixel:

This program is used to find the intensities and absolute positions of individual pixels in a displayed image. Upon calling the program, a flashing cursor will appear in the middle of the TV screen. To find the coordinates and intensity of a given pixel, use the joystick to move the cursor to the pixel and push switch 3,11. The pixel's row and column number will be shown on the L.E.D. display, and with its intensity, be printed out on the TTY. The cursor will then reappear on the screen at its last position. Move the cursor as desired to examine other pixels, or push switch 3,12 to escape the program. Do not use Control-C to escape, however, as it will leave the cursor on the screen. Finally, remember that with compressed images the individual pixels on the screen are really the average of several pixels in the original image. Further, their row and column numbers will always be a multiple of the compression factor.

Box Analysis:

This program is used to examine and numerically analyze a subsection or box of a displayed image. During its execution, the program will access a box you select, and read it into memory. You will then be asked if a printout of the box is desired. If so, the box is printed out as a matrix of pixel intensities labeled with their row and column numbers. (This step can be bypassed by raising the "mute" switch [switch 4,8] beforehand.) The program then prints the center coordinates and size of the box, the location and value of the peak pixel in the box, the total of the pixels' intensities, their mean, standard deviation, and error in the mean.

There are several ways to specify the location and size of the sampling box in the displayed image. Select a box size or request one to be entered from the keyboard by setting switch 2,9 before calling the program. Upon calling the program, a flashing cursor will appear on the TV screen. Select the box's location by moving the cursor to the desired center with the joystick and pushing switch 3,11, or push switch 3,12 to enter the center from the keyboard. Further, by holding either switch down until the bell sounds, the previous box center, and size (if not specified by switch 2,9), will be reused. A final note: All references to box size and location are to the original uncompressed image. If the displayed image is compressed, the size of the box may be modified so that it contains an integral number of final compressed pixels.

Histogram:

This program plots a histogram on the C.R.T. screen showing the intensity distribution of an image's pixels. Histograms can be taken from images stored on 9 track tape, floppy disk, or in the TV display. Set the "data input" switch (switch 4,11) before calling the program, to choose between 9 track tape or floppy disk; you can specify the TV display when the program asks for the image number on the other devices. After calling the program, the histogram is located in the specified image file, or is generated if the file doesn't have one or if the image is held in the TV display. The histogram is then plotted on the screen, showing the relative number of pixels as a function of intensity. For convenience, the mean intensity of the image's pixels is also plotted. Initially, the full range of the histogram over all possible pixel intensities is plotted. You can, however, blow up a section of the histogram for a more detailed examination. A cursor appears on the screen after the initial plotting of the histogram. Move the cursor along the intensity scale with the joystick, and mark off the limits of the intensity region you wish to examine by pushing switch 3,11. The selected region will then be plotted. Push switch 3,12 to replot the original histogram. Hold switch 3,12 till the bell sounds to escape the program. Note: Bypass the image status information printout by raising the "mute" switch (switch 4,8) before calling the program.

Row Examination:

This program can be used to examine and plot out the intensity profiles of individual rows of pixels in images stored on 9 track tape, or floppy disk. Set the "data input" switch (switch 4,11) before calling the program to choose between the two devices. After calling the program, you will be asked for the number of the image you wish to examine, (Raise the "mute" switch [switch 4,8] beforehand to avoid the image status printout.) and a scale factor for plotting the rows on the C.R.T. The program is then ready to display individual rows of the image. Enter the absolute number of the rows you wish to examine. The rows will be found, plotted on the C.R.T., and held in memory. When you have accessed all the rows you wish to examine, or the memory space fills up (which ever comes first), you can plot out the intensity profiles on the Calcomp plotter. The plot will show the rows aligned horizontally by their columns, and positioned vertically by the relative spacing of their row numbers (which are also plotted for identification). Before plotting the rows, however, the program will enter a preview loop which can be used to examine the total ensemble of rows together, and to test various plotting scale factors.

Quick Reduction:

This is a "canned" routine which will reduce raw C.C.D. images by subtracting a "dark" image (scaled by relative exposures times) from them, and dividing the result by a flat field image. This program can only reduce sections of images that will fit completely on one floppy disk. (no more than about 400 by 400 pixels in size) Further, this program works with few user options; it should be used only for a quick "at the telescope" reduction.

To reduce a C.C.D. image, load a flat field image or image subsection into floppy (or dectape) unit "5", and a dark image into unit "6", using the "tape to floppy" program. Set the reduction window beforehand, as desired, to transfer sections of images that will fit on one floppy disk. To hold the reduced image, load a blank floppy disk into unit "7", and write enable the unit. Once these steps are complete, call the reduction program with the "window" switch (switch 4,9) raised beforehand so the preset reduction window is used by the program. The program assumes that the image to be reduced is held on 9 track tape, and asks for its number. The program then asks for the numbers of the flat field and dark images. If desired, the program will skip the subtraction of a dark image, or the use of a flat field image.

The full reduction process begins with reading in a section of the dark and flat field images. The "pedestal" of 40 intensity units is subtracted off from both images. The dark image is then scaled to the exposure time of the flat field image and is subtracted off from it. The corresponding section of the raw image is read in, and has its pedestal of 40 also removed. The dark image is scaled to the exposure time of the raw image and subtracted from it. Lastly, the raw image is divided by the flat image, has a pedestal of 5 intensity units added to it to suppress negative pixel values, and is written to the output disk. This process continues until the entire image is reduced. At the end of the reduction process, the system concludes by generating a histogram record for the reduced image file.

THE SWITCH PANEL

The current assignments of the switches on the switch panel are given below. The uses of the switches should be regarded as provisional, and subject to change as new programs are written and installed in the system.

Switch 4,4: Observation Set-up Programs

0 - Scrub
1 - Focus
2 - Set Telescope
3 - Set C.C.D.
4 - Set Observation Number
5 - Change Header
6 -
7 -

Depress switch 3,2 to call the above programs.

Switch 4,1: Data Input and Output Handling Programs

0 - Save TV
1 - Initialize Tape
2 - Edit Output
3 - Directory
4 - Floppy to Tape
5 - Tape to Floppy
6 - Reduction Window
7 -

Depress switch 3,7 to call the above programs.

Switch 1,10: Image Display Programs

0 - Display
1 - Process
2 - C.R.T. Display
3 -
4 -
5 -
6 -
7 -

Depress switch 4,7 to call the above programs.

Switch 1,4: Data Analysis and Reduction Programs

0 - Pixel
1 - Box Analysis
2 - Histogram
3 - Row Examination
4 - Quick Reduction
5 -
6 -
7 -

Depress switch 4,12 to call the above programs.

Exposure Controlling Switches:

3,3 - Dark	3,4 - Modify Exposure	3,5 - No Display	
3,6 - Pause	3,8 - Start	3,9 - Reset	3,10 - Stop

Exposure time is set with two switches.

Switch 1,7: Time

0 - 1s
1 - 2s
2 - 3s
3 - 4s
4 - 5s
5 - 6s
6 - 7s
7 - 8s

Switch 2,1: Multiplier of Time

0 - 1
1 - 10
2 - 100
3 - Enter Time from TTY

Switch 1,1 is used to specify the observing and reduction windows:

0 - 64 X 64 pixels
1 - 100 X 100 "
2 - 128 X 128 "
3 - 200 X 200 "
4 - 256 X 256 "
5 - 400 X 400 "
6 - 500 X 500 "
7 - Enter size from keyboard

Switch 2,11: Data Output Device

0 - 9 Track Tape
1 - Floppy Disk
2 - TV Display
3 - Quick Look

Program Input Switches:

4,11 - Data Input:	Down for 9 Track Tape	Up for Floppy Disk
4,8 - Mute	4,9 - Invoke Reduction Window	

Switch 2,3: Select Image Processing

0 - Linear Stretch
1 - Log Stretch
2 - Invert
3 - Isophote

Switch 2,9: Select Analysis Box Size

0 - 3 X 3 Pixels
1 - 5 X 5 "
2 - 10 X 10 "
3 - Enter from keyboard

THE SYSTEM MONITOR

The system monitor is the central program of the C.C.D. data taking system, and is responsible for calling in other programs, starting, watching, and reading out exposures, and handling FOCAL or system crashes. The monitor is started at the end of the bootstrap sequence, and simply watches the switch panel waiting for commands, such as starting an exposure or calling in other programs. A few notes on the use of the monitor are given below.

1. When the monitor is first started, it is in a waiting state simply scanning the four program calling switches and the exposure start button. Whenever it enters or returns to this state, a bell will be sounded and the ready light lit. To call a program or start an exposure, set all the other controlling switches beforehand, then depress the appropriate calling switch or start button. The busy light will be lit, and control will be passed to the appropriate program.

2. If a program has been called, control will return to the monitor at its completion. If the program crashes, or is crashed by a Control-C, type "G" (then carriage return) to return to the monitor. Be warned, however, that programs that write to tape lock out Control-C.

3. The monitor can be started directly from FOCAL with the command: X CALL(10,1). It is best to let the system restart itself by typing "G", however, as sometimes, such as with programs using the plotting package, this is the only way the monitor will be properly reset.

4. If an exposure is started, the monitor goes into a different state, watching the exposure control switches in addition to the original program calling switches. No panel lights will be lit while the monitor is watching an exposure.

5. Programs can be called in while an exposure is in progress, although the monitor will be locked out from controlling or watching over the exposure while they execute. A "lockout" light is lit on the switch panel to remind you when the monitor is in such a state. The monitor will regain control over the exposure when the programs complete their execution. If a program crashes, or is crashed, however, the monitor must be restarted by typing "G", as before. If the monitor is not in control of an exposure when it's due to end, the shutter will still close on time, but the image readout will be delayed until the monitor is restarted.

6. Lastly, when control returns to the monitor, a check is always made to see if the "tape write in progress" flag has been set. If it is set, it is assumed that the system crashed while writing to tape (otherwise the flag would have been cleared normally), leaving an open tape file which must be fixed up. The system returns to the monitor state after the fix-up.

THE BOOTSTRAP SEQUENCE

The C.C.D. data taking system is loaded and set up by typing "G" (carriage return), after bootstrapping FOCAL. This starts an automatic "bootstrap sequence", which either prepares the system for initial use and data taking, or recovers it from crashes that occurred previously. The exposure status and observation parameters are preserved over bootstrapping, thus if an exposure was in progress before hand, the bootstrap sequence will be able to reset the system for monitoring or reading out the exposure. Otherwise, the bootstrap sequence will lead you through a series of steps, if desired, to set up the C.C.D. for data taking. A few notes on the sequence:

1. The bootstrap sequence loads the software needed to operate the Cipher tape drive into resident FOCAL. (If a Bright drive is to be used instead, its operating software is already present in resident FOCAL and is not modified.)

2. All steps done in the bootstrap dialog to initialize the C.C.D. for data taking can be done later by calling in individual programs from the monitor.

3. The sequence generates a display on the C.R.T. to announce the system's loading.

4. You can call in the bootstrap sequence from FOCAL, at any time later, with the command: X CALL(10,2).

Below is an outline of the steps taken by the sequence.

1. The first step in setting up the system is to light the busy light, and clear the L.E.D. display.

2. A check is made to see which tape drive is present (if any). If the Cipher drive is present, its operating software is loaded into resident FOCAL. The code needed to operate the Bright tape drive already exists in resident FOCAL, and is not modified if this drive is to be used instead of the Cipher.

3. The "tape write in progress" flag is checked. If it is set, it is assumed that the system crashed while writing to tape. A message will be printed, and you will be asked to confirm that the tape needs to be fixed up. If you concur, a tape fix-up proceeds, and the execution of the sequence resumes.

4. A check is made to see if the C.C.D. controller is connected to the multiplexer. If not, the C.R.T. display is generated, the message "PERFORM DATA REDUCTION ONLY - C.C.D. NOT AVAILABLE" is printed, and control is handed over to the system monitor. Otherwise, the C.C.D. exposure clock is read,

and the sequence continues as below.

5. The C.R.T. display is generated, and the C.C.D. clock read again. If the clock time has changed from the previous reading, it is assumed that an exposure is in progress, and control is handed over to the exposure monitor. Otherwise, the "exposure in progress" flag is checked. If it is set, it is assumed that the exposure ended while the system was down and needs to be read out, as the exposure clock has already counted down. Therefore, the exposure readout is started. Also, if the "readout in progress" flag has been set, you will be asked if the readout should continue. If so, the readout is picked up where it was left off before the system went down; unless the readout was to tape, in which case the the image file in progress will be closed and the output switched over to floppy disk.

6. If this step is reached, it means that the C.C.D. is available for use, but an exposure is not in progress, nor in need of readout. You are then asked, "INITIALIZE C.C.D.? <Y/N>". If no, the observing parameters are displayed and control is handed to the monitor. If yes, then an initialization dialog is followed, as outlined in the "quick overview" section of this report.

WHAT HAPPENS DURING AN EXPOSURE

Below is a step by step account of how an exposure is started, monitored, and read out of the C.C.D. This outline is mainly concerned with the C.C.D. itself; the details of how the data are stored and handled after being read out of the chip are presented elsewhere in this report. An exposure is started by pushing the start button, switch 3,8. The system monitor then proceeds as below:

Starting the exposure:

1. The system responds immediately to the pushing of the start button (provided it's in the monitor state) by lighting the busy light, and clearing the L.E.D. display.

2. The exposure time is computed from switches 2,1 and 1,7. If the time is to be entered from the keyboard, the starting of the exposure is delayed until you do so. Once the exposure time is settled on, it is displayed on the L.E.D. display.

3. The next step is to load the observation parameter list with the exposure time and the "illuminated / dark" status, which is read from switch 3,3. A test is then performed to see which time standards are available, and the U.T. starting time, date, object right ascension, declination, and hour angle are read into the parameter list. If a time standard is not available, however, the starting time is read in from the C.C.D. controller, but the other time and position parameters are not loaded.

4. If the 9 track tape drive is the designated output device, its status is checked now, to avoid a delay of the exposure readout later. A warning is printed and the start sequence continues. If the tape is not ready to store data, you can fix it up any time before the exposure ends.

5. The C.C.D.'s shutter is closed (in case it is accidentally open at this point), the amplifier is set, and the chip is erased twice. The low 12 bits of the exposure time are loaded into the controller, the "exposure in progress" flag is set, and if the exposure is to be dark, a warning message is printed.

6. A command is issued to the C.C.D. controller to start the exposure clock counting, and to open the shutter if the exposure is to be illuminated.

During the exposure:

1. The L.E.D. display is continually updated with the remaining exposure time.

2. Pushing stop (switch 3,10) forces the C.C.D. clock to count to zero, causing the shutter to close, and the readout to begin.

3. The exposure can be aborted by pushing the reset button (switch 3,9). The system will leave the monitor state, but no further action will be taken until you confirm your decision from the keyboard. To abort the exposure, the exposure clock is forced to count out, the exposure flag cleared, and the chip erased twice.

4. The exposure can be paused by raising the pause switch (switch 3,6). During a pause, the clock will be stopped and the shutter closed. The exposure is restarted by lowering the switch.

5. Depressing the "change exposure" switch (switch 3,5) will cause the system to leave the monitor state and ask for a new exposure time. Once the new time is confirmed, the difference between the new time and the elapsed exposure time will be loaded into the controller. Note that the new exposure time refers to the total length of the exposure. If you load a new exposure time less than the elapsed time, the exposure will end.

6. The C.C.D. controller closes the shutter automatically at the end of the exposure. Readout is started when the exposure clock reads zero.

The exposure readout:

1. The first step of the readout is to set the "readout in progress" flag, and to send a "close shutter" command to the controller in the event that the shutter was not closed automatically. The controller's voltages are read, and the output device determined.

2. The observation window parameters are read from the parameter list, unless the "quick look" option has been selected, in which case the whole chip is read out without affecting the parameter list.

3. The selected output device is checked (for details, see the section on that device in this report). The chip amplifier is selected, and the readout begins.

4. The C.C.D. is read out in sets of complete rows, with the chip's shift register being cleared out before each read operation.

5. If the image readout is to 9 track tape or floppy disk, the number of rows read out each time is the largest number of complete rows that will fit into one 4K data buffer in the PDP-8.

6. If the readout is a "quick look", or is to TV memory, as many rows are read out at a time as will fit completely into the top four 4K data buffers of the PDP-8, and still make a set of complete rows in the display after any compression is done. If the image data must be compressed to fit into the display, this is done between each readout.

7. The readout continues until all the rows in the observation window are read out. A running total of the number of rows read out at any point is kept in core, so the readout can be resumed after a crash.

8. At the end of a readout, the output devices are closed (see the output devices sections of this report for details), and if not bypassed, the image and its status are displayed.

C.C.D. IMAGE STORAGE: GENERAL DATA FORMAT AND PARAMETERS

I. Introduction

This section describes the general data storage format and parameters used to record C.C.D. images. Although the particular details on how the data are stored on floppy disks and 9 track tape are different, this section will serve as a good introduction to both of those devices; particularly with regards to the use of the image parameters. For the device-specific details on how the data are stored, see the individual sections of this report on the use of 9 track tape and floppy disks.

II. The Image Format and Parameter List

Each image has a list of parameters associated with it to specify its size, format, C.C.D. condition, and observational information. These parameters are held in the standard 32 word FOCAL I.D. buffer, with the first 31 words used by the C.C.D. system, and the last one used by FOCAL. Most of the parameters are self descriptive, but some explanation is necessary for those describing the image format. This is given below.

Although the C.C.D. itself is of fixed size, you can define an observational or data taking window to make use of an area smaller than or equal to that of the full chip. This window then specifies, by size and location, the actual area of the chip to be used for imaging. The size is specified by a parameter for each dimension, "NR", which gives the image height in rows of pixels, and "NC", which gives the image width in columns of pixels. The location of the window on the chip is also given by two parameters, "SR" and "SC", which give the starting row and column of the image in units of pixels. Both the rows and columns are numbered from zero, with the origin in the upper left corner of the chip, to the number of rows or columns in the chip, less one. With this numbering scheme, "SR" and "SC" also give the number of rows and columns in the chip that had to be skipped over to reach the image origin. Note that "SR" and "SC" will always give the origin of an image in absolute numbers of pixels away from the chip's origin. Lastly, because of the hardware setup of the TV memory, both the number of pixel columns in the image "NC", and the starting column "SC", must be an even number. There is no such constraint on the row parameters "NR" and "SR".

In addition to the four parameters discussed above, two more are needed to completely specify the image format. For some purposes, such as displaying a full image with the limited storage space of the TV memory, it is necessary to average several pixels together. In this process, called boxing, the original image is broken into small cells, each one "B1" rows by "B2" columns, with the pixels in each cell averaged together into one new pixel. The effect of this process, then, is to compress the image by a factor of "B1" in the row dimension, and "B2" in the column dimension. Both "B1" and "B2" are always integers, and are equal to one for uncompressed images. Sometimes the number of image rows, "NR", may not be evenly divisible by "B1"; and the number of image columns, "NC", by "B2". In these cases, left over rows at the bottom of the image, and left over columns at the right end of the rows, will not be used to generate the "boxed" image, and "NR" and "NC" in the output parameter list will be adjusted accordingly. "NC" will also be adjusted, if necessary, to keep the number of compressed columns even. Note, however, with these qualifications, the parameters "NR" and "NC" will always refer to the original image before compression. All of this is outlined in figure 1.

With the above discussion of the image format complete, the full parameter list associated with each image is given below. Basically, the 31 words of the I.D. buffer are divided as follows: Words 1-8, and 22-29 are used to describe the image format and C.C.D. status. Words 0, 9-21, and 30 are used to hold observation information.

Word	
0	The observation number, "NO"
1	The number of data and histogram records, "BF", needed to hold the image data. Note: BF=FITR((NR-1)/RR)+2, where RR=FITR(4096/NC).
2	The number of image rows, "NR".
3	The number of image columns, "NC".
4	The starting row, "SR".
5	The starting column, "SC".
6	The boxing factor for rows, "B1". B1=1 if unboxed.
7	The boxing factor for columns, "B2". B2=1 if unboxed.
8	The Gain + 10*(Bit Select). The bit select is not used now, but will be in the future with 16 bit pixels. It is then used to specify which 12 bits of the full 16 are being stored.
9	Image status, "ST" = 0 For illuminated image. 1 For Dark image. 2 For reduced image.
10 11	Double precision exposure time in seconds.
12	UT start time as 100*Hours + Minutes.
13	Month (UT observation date)
14	Date
15	Year - 1900
16	Hours (Object R.A.)
17	Minutes

FIGURE 1   C.C.D. IMAGE FORMAT

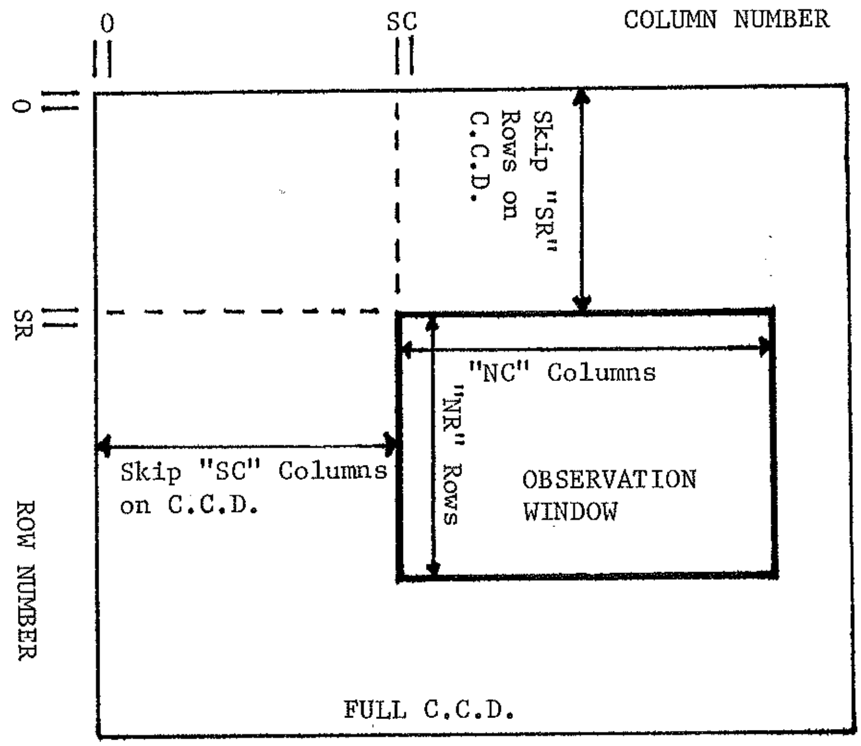

FIGURE 1a   IMAGE COMPRESSION

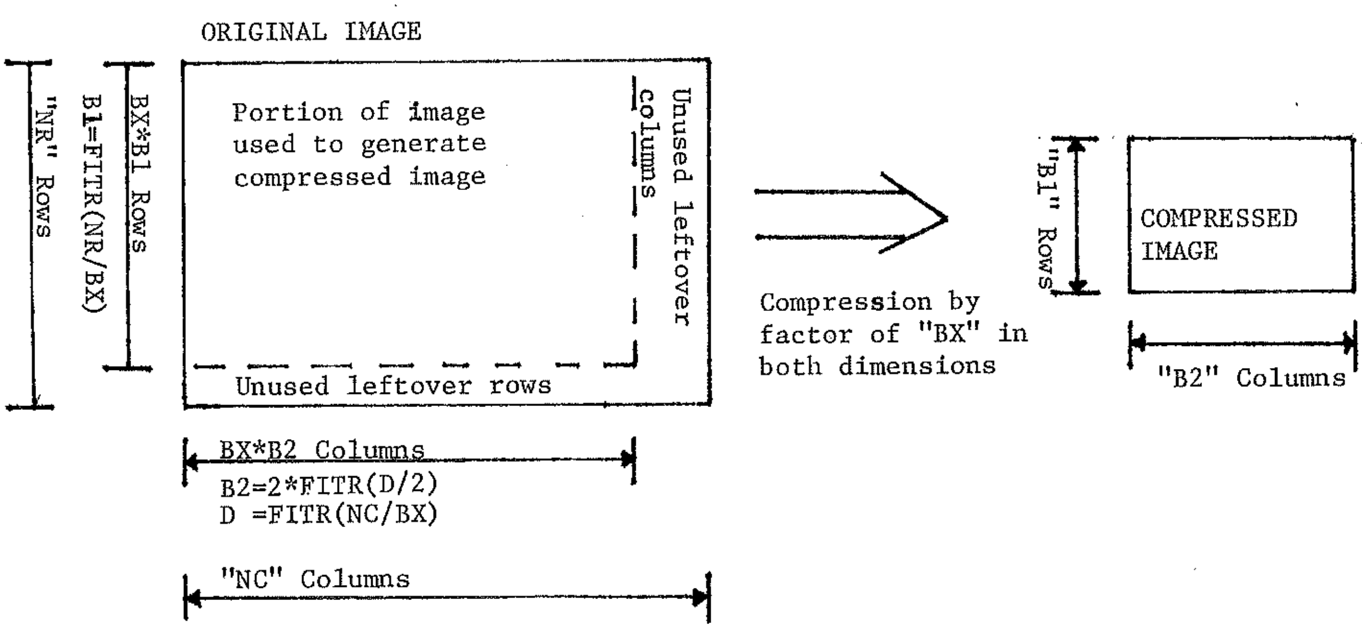

18 Seconds

19 Double precision declination in seconds.
20

21 Absolute value of the hour angle.
HA=100*Hours + Minutes

22 Temperature (C.C.D. voltages)
23 Vertical
24 Horizontal
25 5 volts
26 Drain voltage/3
27 Substrate
28 Unused

29 Amplifier = 0 Lower
1 Upper

30 Image number on tape or floppy disk.

III. Observation Label

In addition to the parameter list, each image has another 32 word I.D. buffer associated with it, which holds an identifying label. This label can hold up to 62 characters of identifying information, object names, or any other type of remarks you desire.

IV. Data Storage

For storage, each image is broken down into a discrete set of chunks or records, each one of which holds a complete subset of the image's rows. Each row consists of "NC" pixels, with each pixel being stored as an unsigned 12 bit integer word. For data read directly out of the C.C.D., a value of "40" is added to each pixel, and a baseline value subtracted from each row. This baseline value is stored as the last pixel in that row. Each record can hold no more than 4096 pixels, so the number of rows stored in each one, "RR", is the largest number of complete rows of pixels which will fit into this limit. "RR" is thus defined as RR=FITR(4096/NC)+1. For the cases when compressed images are stored, it is also necessary to temporarily redefine "NC", as NC=2*FITR(D/2), with D=FITR(NC/B2). From the above, it follows then that the actual number of records required to store the image data is "ND", where ND=FITR((NR-1)/RR). Note that in some cases, the last data record, which will contain the final leftover rows of the image, may contain less than "RR" rows. This will be the case if "NR" is not evenly divisible by "RR". Lastly, if the image is compressed, it may also be necessary to temporarily redefine "NR" as, NR=FITR(NR/B1).

V. The Histogram Record

In addition to the data records defined above, each image has associated with it a full length, 4096 word, histogram record. For raw images read out directly from the C.C.D., this record is blank and contains only zeros, as generating a histogram during data readout is too time consuming. For all other images, however, such as reduced ones or raw images transferred between devices, a histogram is generated and saved in this record.

The histogram describes how the image pixels are distributed in intensity, by giving the numbers of pixels that fall into different intensity bins. Each bin is 5 digital numbers (or intensity units) wide, which means that the histogram will contain a total of 820 bins, beginning with the 0-4 bin, and ending with a separate one for the number of pixels having the maximum value of 4095. Each bin is stored in the record as a double length 24 bit word or channel. The histogram itself begins at word 2048 or channel 1024 in the record. The rounded integer mean of the image pixels is stored in channel 0, and the total number of zero valued pixels is stored in channel 1.

With the above description in mind, it's seen that less than half of the histogram record is actually used. Part of the remaining space will be used to hold further image parameters or information as the need requires. The first segment of it is already used, however, to contain the reduction history of reduced images. The parameters of the original raw image, division image, and subtraction image are stored in their original format beginning at word 100, or channel 50, as follows:

Word	Channel		
100	50	Raw Image:	Parameters
200	100		Label
300	150	Division Image:	Parameters
400	200		Label
500	250	Subtraction Image:	Parameters
600	300		Label

VI. Storage of the Labels and Parameter Lists in the PDP-8

As outlined above, both the object labels and parameter lists are stored in the standard 32 word FOCAL I.D. buffers. The PDP-8 has 6 of these buffers available for use in its memory. The specific uses of the individual buffers by the C.C.D. system are listed below. You can examine any of the I.D. words with the command: S D=FITAK(W,B,L), where "W" is the word, "B" is the buffer, and "L" =0 for single precision, and =1 for double. Likewise, you can set any word to value "D" with the command: X IPUT(W,B,L,D). You should let the system handle the parameter lists for you, however, unless you are trying to test its operation.

Buffer	
-1	Current C.C.D. exposure parameters
0	TV display image parameters
1	Input/Output parameter list handling area
2	Current exposure label
3	TV display image label
4	Scratch and input/output area for labels

VII. Use of the PDP-8 Data Buffers

The PDP-8 has 5 4K buffers available for data storage. Some of the general system uses of the buffers are as follows:

Buffer	
0	Scratch
1	Input/Output data transfer
2	Scratch, and temporary input holding
3	Translation tables
4	Translation tables, plotting package

Be warned that the loading of the plotting package into buffer "4" locks out that buffer from use by FOCAL, until the package is specifically unloaded. While the system takes care of this in all normal usage and crashes, you can do it yourself from FOCAL, if necessary, with the commands: X NAME(12,2);X LONG(-1).

USE OF THE EXPOSURE CONTROL FLAGS, AND CORE VARIABLES

In addition to the parameter list outlined above, the system logic during an exposure and readout of the C.C.D. is also controlled by a set of flags and variables held in the PDP-8's magnetic core. Like the I.D. buffers, the core variables are safe over bootstrapping, and most system crashes. The core variables are also safe over power failures, unlike the I.D. buffers. Each variable is a 12 bit integer word, whose value can be found from FOCAL with the command: S D=FGV(N), and set with the command: X PV(N,D), where "N" is the variable number which is 0-17. The system will take care of all the variables itself, however. You should only change their values for testing purposes, if necessary. How the 18 variables are used is given below:

N

0 The exposure flag. The flag is set to 137 while a C.C.D. exposure is in progress, and 138 when the chip is being read out. The flag is set to 0 otherwise.

1 The pause flag. The variable is set to 1 when an exposure pause is in progress, and is 0 at all other times.

2 A holding area for the number of digits in the exposure time display.

3 The "tape write in progress" flag. The flag is set to 727 when a write to the 9 track tape drive is in progress. It is 0 otherwise.

4 The high 12 bits of the remaining exposure time. This is only used when the exposure is pausing.

5 The low 12 bits of the remaining exposure time, as above.

6 The "readout in progress" flag. Contains a running total of the rows read out so far during the readout sequence. It is set to 0 otherwise.

7 NR (Reduction and display window parameters)
8 NC
9 SR
10 SC

11 NR (Last parameters used for analysis box)
12 NC
13 Center row
14 Center column

15 Number of exposure time extensions. The C.C.D. controller can only hold a 12 bit exposure time. This variable serves as as "lap counter" for exposures longer than 4096 seconds.

16 The number of the current image being written to floppy disk.

17 The number of the current image being written to 9 track tape. Holds the number of the last image written, otherwise. This variable is used to fix up the tape after tape crashes.

USE OF THE 9 TRACK MAGNETIC TAPE

I. Introduction

The nine track Cipher tape drive is the recommended device for the storage of C.C.D. images; indeed at this time, the tape drives are the only devices that can hold a full chip image. It is also recommended that the data be written to tape at high density (1600 BPI) to provide for the highest speed and greatest efficiency of data taking. Used in this mode, a 2400 foot reel of tape can hold about 75 500 by 500 pixel images. For most programs then, one tape should be sufficient for a night's worth of observations. If necessary, the Cipher drive can also write data at low density (800 BPI). This system will also operate the Bright drive, which will only write at low density.

Because of the vagaries of magnetic tape, a considerable amount of engineering has been done to protect you from crashes or errors while using the tape drives. Nevertheless, the tape drives are easy to use. You only have to load the tape, and if writing to it, call the initialization program. The system will handle all tape operations, and should recover from all errors. The main thing to remember is that tape is a sequential storage device. Once you have written a record or file of data, it can not be modified, only deleted, working backwards from the end of the data on tape. The following sections of the report were written to give you a better understanding of how the system uses and handles magnetic tape.

II. The Initialization of Magnetic Tape for Data Storage

When a tape is loaded for data storage, the system has no way of knowing if the tape is blank, has old data that is to be written over, or has old data that is to be added to. The initialization process handles this problem by asking you the tape status. The initialization process then sets up the tape and positions it in the right place to begin data storage. The process also loads the software needed to handle the Cipher tape drive. This is done in the bootstrap sequence, but can be repeated here in the event of severe tape or system crashes if necessary. The system will hang up, and the tape drive will not respond to commands if the controlling software is lost. A few notes on the use of the initialization progress:

1. The initialization process is not needed if a tape is loaded for reading only. Indeed, if it is called and finds the tape without a write ring, it will return to the system monitor without taking any action.

2. If the tape has been moved past the load point, the system will assume that the tape is old. If a write is later attempted, the system will hunt for the end of the data before writing.

3. You should never manually move the tape past the load point. The density burst will not be read, and the tape might be left beyond the end of data. Let the system handle things.

Listed below is the sequence of events used to initialize a tape.

1. The code needed to operate the tape drive is loaded into resident FOCAL.

2. The tape's status is checked. If the drive is not switched on line, a message is printed, and the program interrupted until you switch the drive on line. If the tape is not at the load point, it is rewound. If the tape does not have a write ring, a message is printed, the program stops and control returns to the monitor.

3. If the tape is now ready, the program asks, "WRITE AT END OF OLD TAPE (O), OR START NEW TAPE (N)?"

4. If you designate the tape as old, a search is made for the end of data (which simply consists of two end of file marks), and the tape is left positioned there, ready for writing. If the tape is blank and this option called by mistake, however, several minutes will be wasted while the program searches for a non-existent end of data. Control-C is disabled at this point, but you can stop the program from the PDP-8 console. Push stop, load address, and finally start. Type "G", and restart the initialization program.

5. If you designate the tape as new, you will be asked to confirm your designation before any action is taken which might inadvertently erase old data. On the Bright drives, you will simply be asked to confirm the new tape, yes or no. On the Cipher drives, you will be asked to select high (1600 BPI) or low (800 BPI) for writing. Select high density unless there is a specific need for low. If, however, you have designated the tape as new by mistake, you can escape the program with a Control-C.

6. Once the new tape is confirmed, the program advances the tape past the load point, writing a high density identification burst if necessary. A blank I.D. record of 32 zeros is written, followed by an end of data mark. The tape is left positioned here, ready to write the first image. Control is returned to the system monitor.

III. The Data Storage Format

The general format of image data storage has already been discussed. How this format is specifically used on nine track tapes is presented below.

Tape is a sequential storage device, and therefore has no directory. Each image is assigned a number and accessed by the order of its position on tape. The initialization process writes a dummy image "0" near the beginning of the tape. The first real image on the tape, image "1", is then written after the dummy image. Image "2" follows image "1" on tape, and so on. The image numbers of new images written to tape are computed by reading the image number of the last image on tape prior to the write operation, and incrementing it by one.

All images on tape are treated as separate files. That is, an individual image's data fall between two end of file marks. The first record in an image file is the 32 word parameter list. This is followed by the 32 word label record, and then by the individual image pixel data records. As described before in the "general data format" section, each data record is a complete subset of rows of pixels. The length of each record is exactly that needed to write out its set of rows of pixels, and no more. The last data record, therefore, may be shorter than the preceding ones. The total number of data records in a file is simply that needed to completely contain the image.

After the last data record, the full length 4096 word histogram record is written. This is followed by the last record in the file, which is a repeat of the parameter list. The 31st word in the list is the file's image number, and is accessed to compute the image number of the next image to be written to tape. If an image is the last one on tape, its file is closed out with two end of file marks, which signify the end of data. Otherwise, the image is separated from the next one by a single end of file mark. All of the above details on the data format are shown in figure 2.

Individually, the pixels are stored differently on tape than they are in the PDP-8's memory. Two pixels are written out at a time as a single 32 bit integer. The two pixels are separately the high and low 12 bits of the integer's lower 24 bits. The upper 8 bits are not used for C.C.D. data.

IV. Image Lookup on the Nine Track Tape

The system contains a routine that automatically finds any image file you request. Below is an outline of how it works. The procedure is invoked by all programs that access images on tape.

1. If the tape is positioned at its beginning, the density burst is read, and the tape advanced to the first image.

2. The tape backs up to the start of the current file it's positioned at, and reads its image number.

3. A forward or reverse file offset is computed, and the tape moves to the requested position. If the requested image number is less than one third of the current one, to save time the tape first rewinds, and then starts its search from the beginning.

4. When the tape is positioned at the point where the requested image is believed to be located, that file's image number is read. If the tape is indeed at the correct location, the tape is positioned at the start of the file, and control returns to the calling program. Otherwise, the system tries two more times to find the image. If it is still unsuccessful, the calling program is aborted and control returns to the monitor.

FIGURE 2: DATA STORAGE FORMAT ON 9 TRACK TAPE

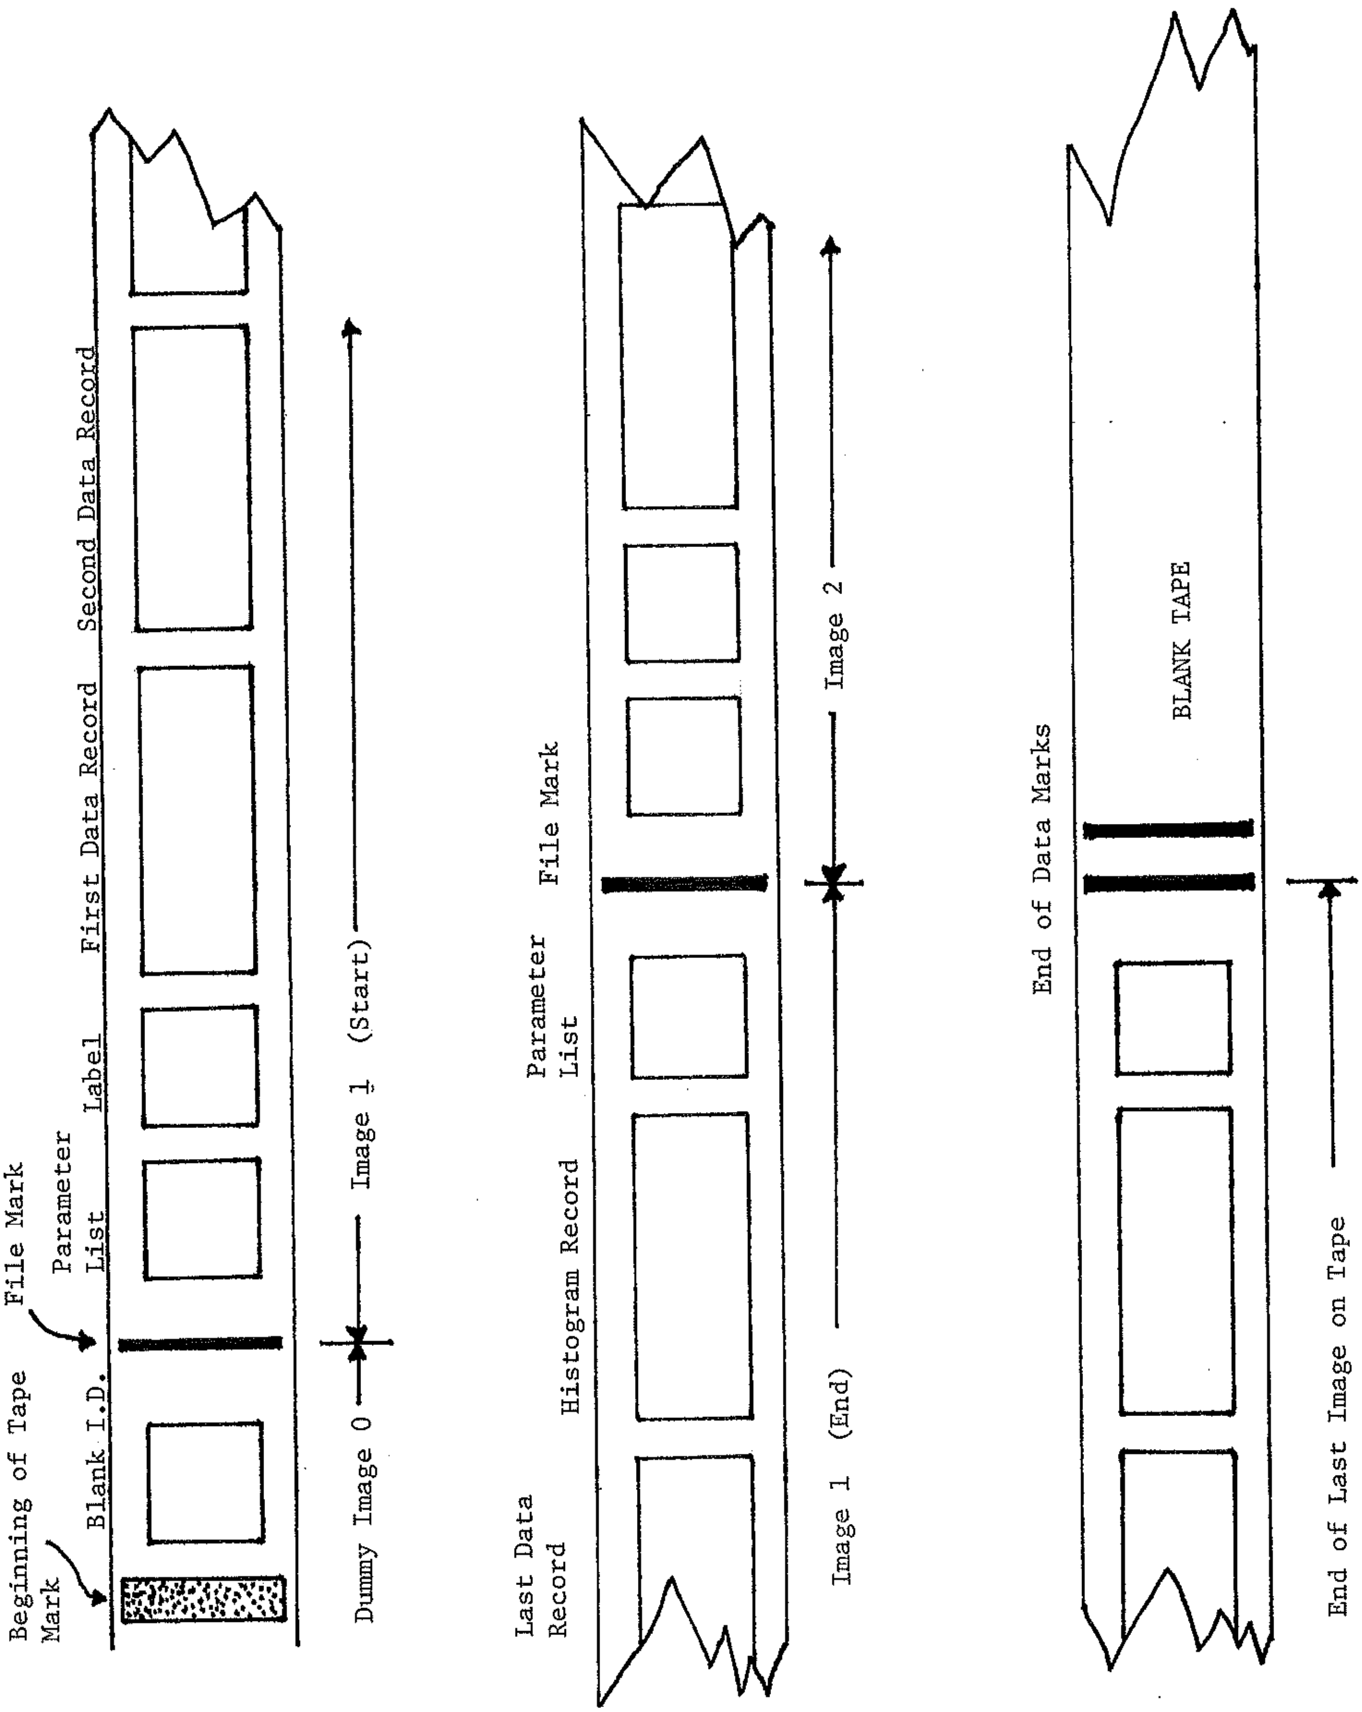

5. If an image number greater than the number of existing files is requested, the tape is left positioned at the end of data, the calling program is aborted, and control returns to the monitor. In practice, this is a good way to manually force the tape to be positioned at the end of data.

V. Tape Status Checks While Writting to Tape

Several checks of the tape status are made before and during all write operations. They are outlined below.

Before the write starts:

1. First the tape drive's line status is checked. If the drive is off line, a message is printed, and the write is delayed until you switch the drive on line.

2. If the tape is found to be positioned at the load point, the initialization program is called. Once you have initialized the tape, the write operation will resume.

3. If the tape drive is not write enabled, is rewinding, or positioned at the end of tape, the write operation will be aborted and control returned to the monitor. If the C.C.D. needs to be read out, however, the system switches the data output over to the backup floppy disk.

Checks while the write is in progress:

1. If a write failure occurs while the system is trying to write a record of data to tape, the write attempt is aborted, and the first part of the bad image file is deleted. If the data were being read out of the C.C.D., however, the file is closed instead, and the remainder of the output is written to the backup floppy disk. A write failure may be due to bad tape, in which case the tape should be unloaded and replaced with a new one.

2. If the end of tape is encountered while writing, the system handles the tape and flow of data as above.

VI. System or Hardware Crashes While Writing to Tape

As described above, the system should take care of all tape errors while a write is in progress. In the unlikely event of a FOCAL or hardware crash while writing to tape, however, the system will no longer be in control, and the tape will be left without end of data markers. Further, depending on the details of the crash, the tape may be left positioned in a blank area, or worse yet, in old data that was to have been written over. To provide for these cases, the system contains a fix-up procedure that will be invoked when you restart or reboot the monitor following a crash while writing to tape. (If a power failure to the tape drive occurs, the drive will stay off line, and will not respond to commands, until you rewind the tape to the beginning and reload it.) This procedure is outlined below.

1. When a write to tape is started, two flags are set in magnetic core, which are safe over bootstrapping and most system crashes. One flag indicates a tape write in progress; the other, the image number of the file being written to.

2. When control is passed to the monitor or bootstrap sequence, the "tape write in progress" flag is checked for. If it is set, it is assumed that a crash has occurred while writing to tape. (In normal operations, this flag is always cleared before control is returned to the monitor.) If the flag is detected during the bootstrap sequence, you will be asked to confirm the tape crash in case the flag was set prior to the initial use of the system.

3. If a crash has occurred, the system locks out Control-C and attempts to find the last good image file. Usually it does this by rewinding to the beginning of the tape and counting forward. To save time during a data taking run with the C.C.D., however, the system simply backs the tape to the last good image file, and checks its date of observation against the current one. If the dates match, the system assumes that it's positioned in new data, and proceeds directly to the next step. Otherwise, the tape rewinds and the above procedure is followed.

4. Once the tape is positioned at the last good image, the tape is moved to the end of the file, and an end of data mark is written. This effectively deletes the file being written to when the crash occurred. The tape is left positioned at the end of data, the flags are cleared, you are informed of the tape's status, and control returns to the monitor.

In the extreme event that this automatic procedure does not work, you can try the following fix-up.

1. Escape from the monitor with a Control-C. If Control-C is disabled, from the PDP-8's console push stop, load address, and start.

2. Control is now in FOCAL. Enter the command X PV(3,0) (always follow all commands with a carriage return).

3. Type "G" to restart the monitor, and request a tape directory. The directory will list out all the images taken so far. If the tape reaches a point where it continues to advance, but does not continue the directory listing, it may be assumed that the last image listed is bad.

4. Stop the directory program or monitor with a Control-C, and rewind the tape with the command X RWND(1). Once the tape has rewound, enter the commands X PV(3,727), and X PV(17,IM), where "IM" is the number of the bad image.

5. Type "G" to restart the monitor. The automatic fix-up should now proceed correctly. Control will return to the monitor when it is complete.

USE OF FLOPPY DISKS OR DECTAPE

I. Introduction

Floppy disks and dectapes are secondary data storage media used as backups for the 9 track tape drive and temporary output for the reduction programs. Unfortunately, they can not even hold one full C.C.D. image. The largest sized array that can be stored on a single disk is slightly larger than 400 by 400 pixels. Nevertheless, full C.C.D. images can be stored by spanning two or more disks. For most operations, the default floppy unit number is "7", although units "5" and "6" are used in the reduction process. It is recommended that you use floppy disks rather than dectapes, because of dectape's extremely slow access speed.

II. The Data Storage Format

The general image storage format has been described earlier. This general format is implemented on floppy disks in a way to provide the greatest possible match with the data storage format used on the 9 track tapes. As on the tapes, images are numbered by the order in which they were stored, although the disks have no dummy image "0" as the 9 track tapes do. Only complete image files are stored on disk. As described below, if a single image is too big to fit on one disk, it is broken up into smaller selfcontained files that will.

A floppy disk can be considered to be composed of 46 separate 4096 word records, each one of which has an associated 32 word I.D. record. For the greatest speed of operation, image data are written to the disk with the FOCAL "MSAV" command, which writes to any of the 46 data records and their associated I.D. records in one operation. An image file on disk can thus be viewed as a collection of an integral number of the data and I.D. records.

The first part of a floppy image file consists of the pixel data and the parameter list. As described earlier, the image data are broken into complete subsets of rows of pixels. Each subset is written to a separate record on the disk, in sequential order, until the image is stored completely. Because the data records have a fixed 4096 word length, there may be leftover words at their ends which will contain meaningless data if the subset of image rows stored in the record contain less than 4096 pixels. A copy of the parameter list is stored in each of the I.D. records associated with the data records. The image file is closed with the histogram record, which follows the data records. The image label is held in the histogram's I.D. record. Note that regardless of how small an image may be, a minimum of two floppy disk records will be required to hold it; one for the data, and one for the histogram.

III. Image Lookup

Floppy disks, like the 9 track tapes, do not have a file directory. When you request the system to lookup an image on disk, it accesses the second word in the parameter list of the first image file on disk. This word tells the system how many records are contained in the first file, and hence how many records to jump across to reach the second file. The system continues to hop from file to file, by accessing each file's "record count" word, until the requested image is found. If the record count word is zero, or specifies a greater number of records than can be stored in the remaining disk space, it is interpreted as an end of data mark.

IV. The Initialization of Floppy Disks

Disks that have a zero in the "record count" word in the first I.D. record on the disk, are properly initialized for the writing of their first image. The writing of the first image file concludes with the zeroing of this word in the first I.D. area of the next image file, and so on. The zero is interpreted as an end of data mark, and hence marks the starting point for the next file to be written. Blank newly formatted disks contain zeros, and do not need to be initialized; likewise old disks that are to have new files added after the old ones. Old disks that are to be written over should be initialized, however, by calling the edit program (with the input switch set to floppy), and clearing the disk. The clearing operation simply writes the initial zero in the first "record count" word.

V. Disk Status Checks and Crashes While Writing

Checks of the disk status are made before and during write operations. They are outlined below.

Before the write starts:

1. The end of data is located on the disk, and the amount of storage space remaining is computed. If the image will fit on the disk, the write operation proceeds. If the image will not fit in the remaining disk space, but would fit on a blank disk, the write is delayed until you load a blank disk. If the image is too big to fit on even a blank disk, the file is broken into pieces as described below, a warning message is printed, and the write proceeds.

Checks while the write is in progress:

1. If a write error occurs, the program is interrupted until you take action, such as checking the loading of the disk or write enabling the unit. If the disk is bad, it can be replaced, and the write will continue. The new disk will not have the initial records, but can be patched up later. (At this time, there is little use of the floppy disks for primary image storage and the "patch-up" program has not been written yet.)

2. If the disk fills up, it is assumed that image spanning will occur. The old disk is properly closed, and the write operation is interrupted until you load a new disk.

3. At the end of the write operation, the system sends a warning message if the remaining space on the disk is insufficient to hold another image of the same size as the one just written.

VI. Image Spanning

If an image file is too big to fit completely on one disk, it is broken into smaller image files that will. Each small file will consist of a complete subset of the image's rows, with the "starting row" and "number of rows" parameters adjusted to indicate which rows are stored in which files. These files will then be treated as separate images on the disk. (At this time, a program to splice a spanned image back together has not been written yet.)

When taking data from the C.C.D., it is recommended that you use blank disks for the image spanning to minimize the number of disks used, and to prevent mixing in other images on the same disk. Further, the automatic image display after the readout will be disabled if spanning occurs.

USE OF THE VIDEO DISPLAY

I. Introduction

The primary display device for C.C.D. images is the area scanner memory coupled to a standard video monitor. For display purposes, the memory can be thought of as a 256 by 256 array of pixels. The memory displays the array as a standard TV frame, with each pixel assigned one of 64 possible gray levels. The gray levels are generated from a 6 bit slice of the 12 bit pixels, which you select according to which features of the image you are interested in examining.

Images can be called in for display from the 9 track tape, floppy disk, or directly from the C.C.D. as a way of quickly examining a field of view with out taking the time to save it on tape first. Indeed, the "quick look" option of data taking has been provided for this purpose, and will read out the full chip directly to the display without affecting previous observing parameters. Images held in the display memory can be dumped out to tape or disk later for permanent storage. The display memory also provides you with a convenient way to interact with the images. Several programs are available in the system that will enable you to directly examine individual pixels, perform statistics on selected areas of the images, or process them to highlight different features in the objects under study. The main limitation in the use of the display memory is its small size. Images larger than 256 by 256 pixels in either dimension must be compressed and therefore degraded in resolution before they can be displayed.

II. Use of the Video Display

To set up the video display for use, the area scanner memory must be turned on, and the stop and reset buttons pushed. (As of this writing, you must also provide a TV monitor, which should be patched into the video out terminal of the area scanner memory.) The memory must also be write enabled; the system will print a warning message if it is not.

To actually display C.C.D. images, call the display program and specify the image you are interested in. The system will find the image, and center it in the monitor's screen. The actual appearance of the display is controlled by three things, the contrast and brightness settings of the video monitor, and the bit select on the memory. The bit select is set by a thumbwheel on the front panel of the memory, and controls which 6 bit slice of the 12 bit pixels is to be used to generate the gray levels for the display. A bit select of "0" will display the lowest six bits of the pixels, and can be used to examine faint features or subtle gradients of surface brightness in the image. Note that a bit select of "0" will also break the image down into a set of isophotes with a range of 64 digital numbers (out of the full 4096) between each isophote, with each gray level between isophotes representing an increment of one. A bit select of "1" will show the next highest slice of six bits, generate isophotes with a spacing of 128 digital numbers, grey levels with an increment of 2, and so on, up to a bit select of "6" which will show the highest 6 bits of the pixels, not generate any isophotes, but will have each gray level equal to an increment of 64. Further, a bit select of "7" will show the highest 5 bits as

FIGURE 3: THE VIDEO DISPLAY FORMAT

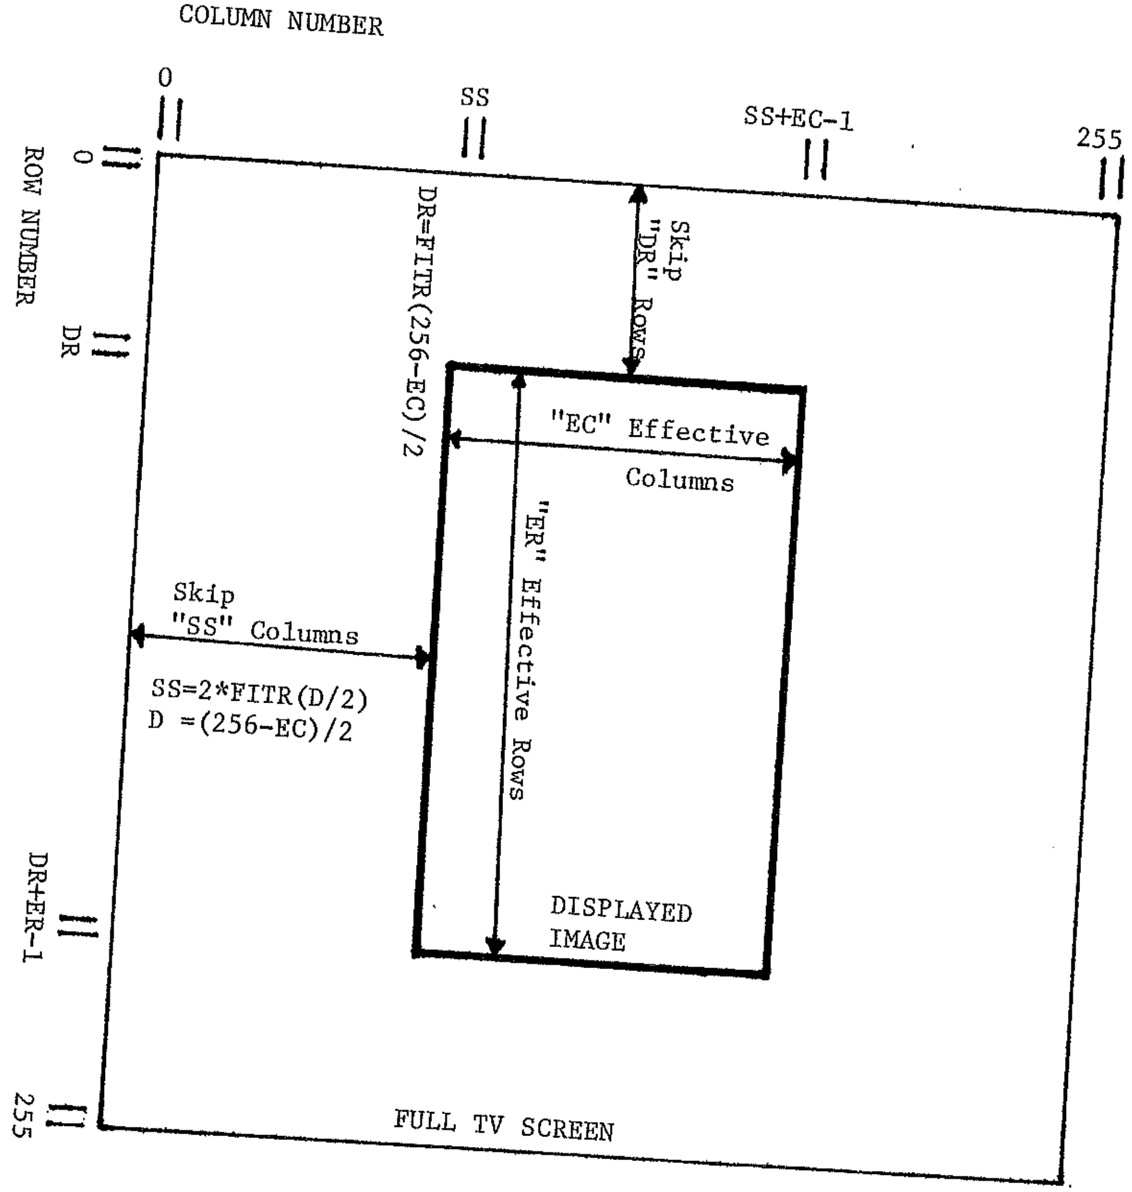

the lower 32 gray levels on the display. The bit selects "8" to "15" will simply repeat the above, with "8" equivalent to "0". Finally, be warned that the quality of the display may be affected by the shape of the image raster. The raster is often stretched out in the horizontal direction, and can cause such things as the spurious appearance of oval star images.

III. How the Images are Written to the Video Display Memory

Storage space in the memory is addressed by row and column numbers, both of which range from 0 to 255. The displayed image can have any number of rows within these limits, but each row must contain an even number of columns. The rows must also start on an even column boundary in the memory. This is because the memory hardware is designed to hold 24 bit words, hence pixels can only be written to the memory in pairs. Although this problem could be avoided with the appropriate software, in practice the necessary manipulations required to display images with an odd number of columns would take too long. The main problem with the display, however, is to provide for a way of displaying images larger than 256 pixels in either dimension. How the images are handled is outlined below, and in figure 3.

1. The first step is to compute an image's effective number of rows and columns, "ER" and "EC". These are computed from the image's original dimensions and compression factors as: ER=FITR(NR/B1), and EC=FITR(NC/B2). Usually "B1" and "B2" are one, and the image is uncompressed. The procedure of finding the effective number of rows and columns, however, covers the cases when an image is stored in an already compressed format.

2. If both "ER" and "ER" are less than 257, the image can be read directly into the display. Otherwise, the image must be compressed by a factor of "BX" in both dimensions (to preserve the aspect ratio) so it will fit. "BX" is computed as BX=FITR((N-1)/256)+1, where "N" is the larger of "ER" and "EC".

3. If the image is to be compressed, its effective dimensions are changed from "ER" and "EC" to "BR" and "BC", where BR=FITR(ER/BX), and BC=2*FITR(D/2), with D=FITR(EC/BX). The compression is done by starting at the image's origin, and dividing it into cells "BX" by "BX" pixels across. The pixels in each cell are then averaged together into one new pixel, making the final displayed image an array of these averaged pixels. Note that the last leftover "ER-BX*BR" rows at the bottom of the image, and the last leftover "EC-BX*BC" columns at the ends of the rows, will not be displayed.

4. The final displayed image will be centered in the monitor for ease of viewing. This means that the image actually begins in row "DR" and column "SS" of the memory, where DR=FITR((256-ER)/2), and SS=2*FITR(D/2), with D=FITR((256-EC)/2).

5. The image's parameter list and label are stored in two of the PDP-8's I.D. buffers, with the parameters adjusted according to the final number of rows and columns actually displayed, and the final compression factors used.