

WORKING PAPER 70

GREY SCALE DISPLAY SLAVE

by

MICHAEL BEELER

Massachusetts Institute of Technology

Artificial Intelligence Laboratory

May, 1974

Abstract

The Programs SNAP and ZSLAVE are components of a new grey scale display system. The object is to produce photographs, from a computer display, which have grey scale resolution comparable to that of the visual input devices and the vision data at the A.I. Lab.

Work reported herein was conducted at the Artificial Intelligence Laboratory, a Massachusetts Institute of Technology research program supported in part by the Advanced Research Projects Agency of the Department of Defense and monitored by the Office of Naval Research under Contract Number N00014-70-A-0362-0005.

Working Papers are informal papers intended for internal use.

Grey Scale Display Slave

The programs SNAP and ZSLAVE are components of a new grey scale display slave system. The object is to produce photographs, from a computer display, which have grey scale resolution comparable to that of the visual input devices and the vision data files at the A. I. Lab.

The program SNAP runs as a job under the time sharing system. SNAP asks the system to assign the PDP-6 to it, and then SNAP loads the PDP-6 with ZSLAVE. The user communicates his desires to SNAP, which then reads in the appropriate file of vision data, converts format as necessary, and loads point or line display commands into a buffer area in the PDP-6. ZSLAVE reads from this buffer and displays the required points, either on the DEC type 340 display or on the Tektronix D10 display, connected to digital-to-analog multiplexor channels.

using SNAP

When SNAP is started, it may, under unusual conditions, halt at mnemonically named locations. These are:

NOPDP6 -- time sharing system would not assign PDP-6 to SNAP;
probably someone else has it
NOCORE -- time sharing system would not map PDP-6's core
into SNAP's core image
NOSLAV -- no ZSLAVE file was found for loading into PDP-6
NOLOAD -- time sharing system would not load ZSLAVE file
into PDP-6

A more common error condition is that the PDP-6 does not respond; in this case an informative message is typed:
PDP-6 DEAD? RESTART WITH IT EXECUTING 41.

Assuming none of the above errors occur, the message,
TYPE IMAGE FILE NAME
is printed. A standard ITS file name specification is expected, terminated by carriage return. If the file is not found, SNAP asks for a file name again.

SNAP reads in the file and looks at the first character. If it is "(" (open parenthesis), the file is assumed to be in Lisp format. Otherwise, FAKE TV format is assumed (see section below). Lisp format files consist of the following characters (in ASCII, packed left-justified, 5 to a word):

(and)
0 through 9
minus sign
space

. and carriage return and line feed -- these 3 are ignored
The digits 0 - 9 make up decimal integers. A minus sign anywhere in a number negates the number.

The parentheses must follow the following convention:
the vision data file is:

((block #1)(block #2)...(block #m))garbage

where "blocks" are:

X Y Z

or X0 Y0 dX dY Z1 Z2 ... Zn

An X Y Z block specifies one point at coordinates (X,Y) and with brightness Z. An X0 Y0 dX dY Z1 Z2 ... Zn block specifies a line of points, the first at (X0,Y0) and with brightness Z1, the second at (X0+dX,Y0+dY) and with brightness Z2, etc.

Each number is terminated by a space. The numbers obey the following constraints:

$0 \leq X, Y, X0, Y0 \leq 16383$ (decimal, = 37777 octal)

all Zs must be non-negative and there must be at least one Zn

Within the blocks, arbitrary placement and depth of balanced parentheses is accepted, so both

(X0 Y0 dX dY) Z1 ... Zn

and X0 Y0 dX dY (Z1 ... Zn)

are valid, for instance. "Garbage" following the close parenthesis which matches the initial open parenthesis is ignored.

Before displaying begins, the SNAP program waits for any one of the spacewar colsole switches to be pushed, or for a space to be typed. This permits the user to synchronize displaying with the manual operation of the camera shutter.

FAKE TV format

In the case of a Lisp-format file, described above, the SNAP program prints the message, "FIRST CHAR IS (SO LISPY FILE ASSUMED," and starts to work. With FAKE TV format, however, several further user options are available. This is because FAKE TV files are essentially raw vision device data, and various "massaging" of the data may be appropriate. For a complete description of the FAKE TV file format, see vision group documentation. A brief summary follows.

The first block contains the coordinates of the lower left and upper right corners of the image (0 to 16383 is full scale). The dim cutoff level (DCL) and confidence level (CFL) at which the image was read in are next, optionally followed by a descriptive title enclosed in parentheses. The rest of the first 1024 words are presently unused.

So far, all vision files have been recorded at a linear scanning density of 1/16 that of the 0 - 16383 grid size, which is 1/256 of all possible points in the image. The image is scanned in square blocks 64 measured points (= 64 * 15 grid points) on a side, creating 64 * 64 measured data points per block. This data is packed 4 points to a computer word, yielding 64 * 64 / 4 = 1024 computer words per block. The SNAP program follows the vision group convention of using assembly parameters for these quantities, and may need changing if FAKE TV format parameters are changed. Points within blocks and blocks within the image are scanned to the right along the bottom row, then the next row up, etc.

using FAKE TV options in SNAP

When processing a FAKE TV file, SNAP first types, "FAKE TV; SAME?"
If the user merely wishes to re-process the current file with options unchanged, typing "Y" accomplishes this. Otherwise, informative verbiage appears:
FIRST CHAR IS NOT (SO FAKE TV FILE ASSUMED
XLL = 6400. YLL = 4000. XUR = 12528. YUR = 11152.
DCL = 1 CFL = 3
(COMMENT DUMP OF A PRINTED CIRCUIT BOARD UNRETOUCHED PHOTO
MAY 18 1973 FININ)

The user is then queried about options:

CENTER & MAGNIFY IMAGE TO FILL SCREEN?

If "Y" is typed, the X and Y coordinates will be scaled up by the (same) integer which is as large as possible without making the range of either exceed 16384. X and Y are also offset as needed to center the image on the display screen. Then, or if something other than "Y" was typed:

SUBTRACT DATA FROM 1300? Y, N, OR OCTAL NUMBER

This allows vision device data, which may represent amount of darkness rather than amount of brightness, to be inverted. The default (obtained by typing "Y") value of 1300 is a typical maximum value returned for blackness by the vidisector "TVC." The user may wish to use a lower value for the inversion, especially if his data's maximum value is, for example, only 670. Numerical type-in is terminated by carriage return. Rubout is implemented.

DIVIDE RESULT BY WHAT (OCTAL)?

Typing 1 (terminated by carriage return) takes the data as is. The resolution of the data may, however, exceed the resolution with which the user wishes to display it. This is especially true when making test runs, since non-scaled data could require a minute or more for completion.

The SNAP program now is ready to process the data, and depressing any spacewar console switch, or typing a space, will cause it to proceed with the displaying.

The FAKE TV routines in the SNAP program check Z values for being non-negative. Failure results in the message:

NEGATIVE Z DATA -- ABORTED

Abortion can also be accomplished by typing "A" while the SNAP program is processing either a Lisp or a FAKE TV format image.

When the SNAP program finishes loading image data into the PDP-6, it types "LOADING DONE," and when it detects that the ZSLAVE program in the PDP-6 has finished displaying, SNAP types a bell and "DISPLAY DONE." A new image file specification is then requested, etc.

One final note on photographs of FAKE TV format data:

Faintly darker lines seem to appear between the blocks of 64 X 64 points. These lines are not entirely explained by deflection hum (see below), and may be due to phosphor emission or intermittency effects (also see below).

the ZSLAVE program

The slave program run in the PDP-6, ZSLAVE, interprets commands placed in a ring buffer in PDP-6 core memory and intensifies points on a display screen accordingly. The data switches on the PDP-6 console specify which display to use. If bit 0 (leftmost) is on (up), the DEC type 340 display is used. If bit 0 is off, the Tektronix D10 display is used via digital-to-analog multiplexor channels. Bit 35 (rightmost) being on selects a test pattern to allow adjustment of the D10's intensity control knob. We will return to discuss these displays after finishing with the slave program itself.

The ZSLAVE program reads commands out of a ring buffer, thus allowing the SNAP time-shared job and the ZSLAVE display routines to run in very loose synchronization. In practice, there are occasions when one is waiting for the other, and also occasions when the situation is reversed.

The data in the ring buffer is structured to achieve interlocking between the two programs, and also to distinguish between a single point command (X Y Z) and a line command (X0 Y0 dX dY Z1 Z2 ... Zn). The data structure is documented in the ZSLAVE source file. Its conventions will allow other routines, such as LAP code similar to the current Lisp display slave or the Lisp color scope display slave, to run the ZSLAVE program easily.

philosophy and photography

The emulsion (light-sensitive layer) of photographic film contains silver halide crystals. When struck by enough light, a portion of such a crystal dissociates -- some halogen atoms leave their places in the crystal structure. The development process in a darkroom turns any such crystals totally into granules of metallic silver, which is so fine that it looks black. Intact silver halide crystals are washed away, and a photographic negative results.

There are many variables in determining how many silver halide crystals get developed. Film speed (sensitivity of the type of film), diaphragm opening ("f stop"), and incident light intensity are the main factors. (A guide to proper exposure settings, written by Berthold Horn, is included as an appendix.) Shutter speed, normally a factor, is often irrelevant in photography of computer displays. This is because the shutter is usually opened prior to any displaying and held open until the image has been completed. In place of shutter speed, then, we have intensification duration -- how long each point in the image is actually emitting light.

Unfortunately, this duration is determined not only by the amount of time the display tube's electron beam is on (which is typically well controlled), but also by the intensity of the electron beam (which is typically repeatable but non-linear; see the appendix), and the light emission characteristics of the display tube's phosphor. The phosphor may emit different colors at different times after electron bombardment, such as the P7 phosphor on the DEC type 340, which has bright blue, rapidly decaying emission and greenish, slowly decaying emission. The photographic film used may be noticeably more sensitive to one (usually the blue) than the other. The Kratos color scope's phosphor emits reds or yellows or greens, depending on the energy with which electrons impinge upon it.

High electron energy is used in many displays to achieve high brightness. As the energy is increased, however, secondary emission (of electrons knocked out of the phosphor by the energetic beam) may create a halo around displayed points. This halo is usually faint, but could conceivably cause fuzziness around edges in the image.

For all these reasons, computer displays are usually adjusted and photographed as follows:

(1) Usually the display is already adjusted so that appreciable halo from secondary emission is absent.

(2) Usually the energy given each electron in the beam, determined by the anode voltage, is not variable. (The computer has control only over how many electrons are sent to the screen.)

(3) Often the photographic film, or optical filters mounted on the camera, are carefully chosen to match the phosphor's light emission characteristics. (This is not very crucial in black and white photography.)

(4) Computer hardware usually intensifies (turns on the electron beam with it directed at) each point in the image for a fixed amount of time. Variable intensity is achieved by selecting any one of a few (typically eight) fixed values of electron flow rate. These values are monotonic, but non-linear, and the size of the spot typically varies from one value to another.

For the purposes at hand, the conventions in (4) above are not acceptable. We want several (say 64 or more) intensity levels, not just eight. We would like to have them linear and of constant spot size. The solution adopted is to re-display the same point several times; the number of repetitions governs the intensity. The value of electron flow rate is not changed.

Unfortunately, this does not guarantee complete linearity. Phosphors may not emit twice the total amount of light when bombarded by two closely timed pulses instead of one pulse. (The photographic film may also introduce non-linearities; see below.) Additionally, the spot size almost always increases with intensity. Since the electron beam's boundary is not well defined, but more like a Gaussian distribution, several repetitions of the same Gaussian will add up to a larger diameter of spot above exposure threshold. This is unavoidable, and we must simply select a display whose electron optics are capable of focusing the beam well and whose total display area is large relative to spot size.

Another factor, which we shall blithely ignore, is possible non-uniformities in the phosphor coating on the screen of the display tube.

CONCLUSION: The Z values given to ZSLAVE are the number of times to repetitively display each point.

quirks and technicalities

reciprocity failure:

A photographic emulsion responds differently to the same total amount of light, depending on whether the exposure (duration of the light) is short or long. Usually a very short exposure reduces speed (sensitivity of the film to light) and contrast -- and, especially with less modern films, more for fast films than slow. Thus with exposures of 100 microseconds or less, a medium or slow speed film may be as fast as a higher speed one; and the high speed one, especially, may have reduced contrast -- even though the total amount of light is the same as for a longer exposure. (This may be due to the time it takes electrons to migrate around in the silver halide crystal, and for development spots or silver atoms to be formed from the halide. High speed films generally have larger crystals than slow films.)

Very long exposures also lose speed. Kodachrome exposures over 30 seconds or so, for instance, don't have nearly as much effect as shorter exposures at correspondingly higher intensity.

Reciprocity failure is probably irrelevant to photography of computer displays, but is included here for the sake of completeness. Although the amount of time the electron beam is on may indeed be only a few microseconds, the phosphor probably emits light for a moderately long period of time. (Not terribly long or light pens wouldn't work!) Also, the dimmer-but-slower green emission of P7 phosphor may integrate up to a total light contribution which is not negligible. The long exposure effect is irrelevant because, although the shutter may well be open for a minute or more, the light is actually present at any given point for a more limited time.

In truth, the phosphor's emission response and decay properties could be important to a user attempting to create carefully controlled density levels in his photograph. Besides trial and error, I suggest borrowing a photodiode from MIT's Strobe Lab. Such units plug directly into an oscilloscope and are pre-calibrated, so the true response/decay can be easily measured.

intermittency effect:

A series of short exposures don't have as much effect as a continuous exposure of the same total amount of light. It's possible to get exposures so short nothing happens, no matter how many times it's repeated. (Perhaps due to electrons snapping back to where they were when the photons knocked them loose.) This effect may be thought of as an extreme case of reciprocity failure. Again, it is probably irrelevant in the case at hand.

display characteristics -- DEC type 340

Of the sundry ways to re-display the same point on the 340, the following has been empirically determined to be fastest:

- (1) deflect to the desired point and enter increment mode
- (2) send pairs of increment mode halfwords which say:
 move +x & don't intensify, null, null, null
 move -x & intensify, null, null, null
- (3) repeat step (2) as needed by BLK0ing

This process runs at about 20.6 microseconds per pair of halfwords. Since the electron beam is on for only 1.5 microseconds or so, efficiency is very low. The brightest hardware intensity level (7) is used to avoid further slowness. If the desired point is at the right edge of the screen, the table of increment halfwords must have the signs of its move +x and move -x inverted.

display characteristics -- Tektronix D10

The Tektronix D10 has X, Y, and Z (intensity) inputs of sensitivity commensurate with the digital-to-analog converter output multiplexor. (0 to 37777 octal yields about -10 to +10 volts. X and Y receive this full range; only -1.25 (off) and +1.25 (on) volts are currently used on Z.)

Any other suitable oscilloscope may be used in its place. Currently, DAC channels 15, 17, and 14 are X, Y, and Z, respectively. (One should check cross-talk between channels, linearity, and crossover distortion before changing these assignments.) With PDP-6 data switch 35 on, two dots are displayed, both at Y = 0 volts. At X = -5 volts a dot is displayed with Z = off, and at X = +5 volts a dot with Z = on. The user should adjust the D10's intensity control knob so the former is not visible and the latter is as desired.

Since the D10's beam is continuously on while a point is being displayed, efficiency is high; the same image can be presented much faster than with the 340 display. When using the D10, the Z value given is how many SOJGs to execute in an AC, which take about 2.5 microseconds each. (For Z > 0, there is actually an overhead of about one SOJG worth, during which time the beam is on. This is generally negligible.)

shortcomings in displays -- DEC type 340

Currently the 340's beam-off bias is adjusted so that the spot is faintly visible. While waiting for more data, the ZSLAVE program re-positions the beam to (0,0), thus avoiding creating bright dots at random places in the image.

Currently the 340 suffers from an intermittency in Y deflection which causes each portion of the image to jitter toward the top of the screen by perhaps 5 percent of its distance thereto. It also suffers from 60 cycle hum, particularly in its Y deflection, where the hum is at least one deflection unit in amplitude. This hum has the effect of superimposing Moire patterns on the desired image. The result, looking like a painting on wood whose grain shows through, is beautiful, but...

Another, more fundamental difficulty with the 340 is its curved screen. This causes distortion of the image as it is focused onto flat film. A square outlining the addressable area on the 340, when photographed, has the middle of each edge bowed in toward the center. This type of distortion is called pincushion distortion, as opposed to bowing outward from the center, which is barrel distortion.

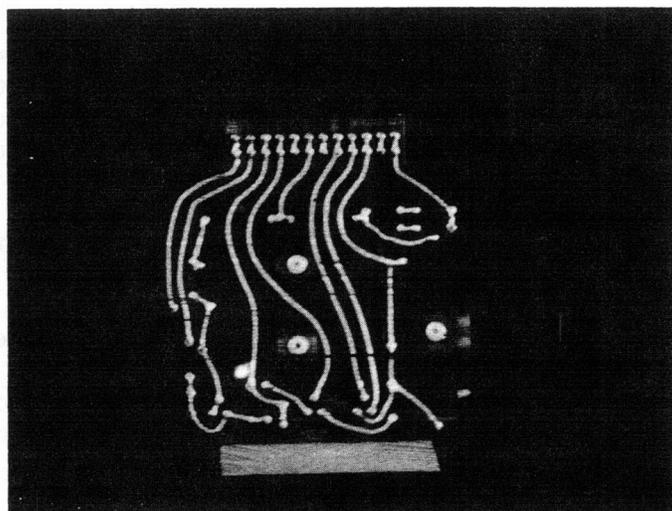
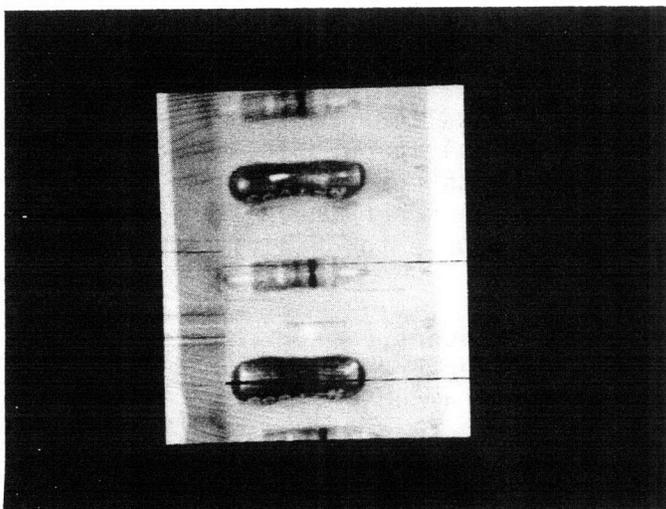
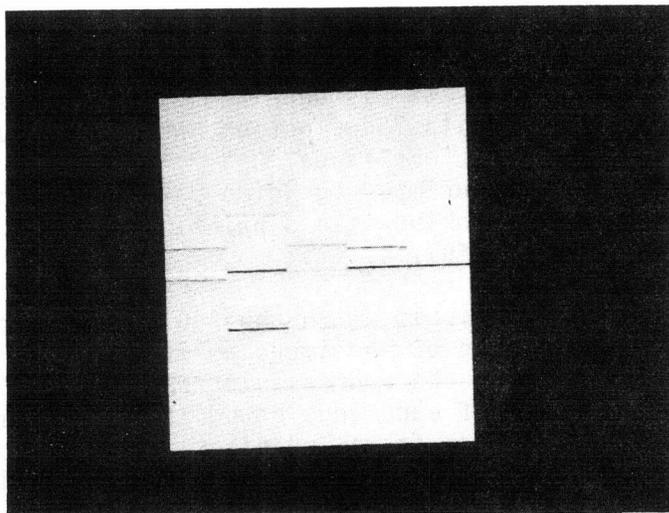
shortcomings in displays -- D10

Note that the D10's ruled graticule is removable, exposing a flat screen. Thus optical distortion encountered with the 340's curved screen is not present. The electron optics in the D10, however, cause the image on its screen to suffer slightly noticeable barrel distortion.

The D10 also suffers from hum, of approximately equal amplitude on X and Y. Unfortunately, Tektronix admits that a small amount of deflection hum is one of the idiosyncracies of the 5103N family of oscilloscopes, of which the D10 is a member. It is also unfortunate that home-brew input circuitry has been installed (partly by me, and partly for scrapped vidisector monitor plans) in our D10 instead of purchasing the \$25 5A24N buffer amplifier with centering control. Not only will it now be hard to obtain Tektronix service, but there is no mechanism for centering the displayed image to obtain full use of the screen area.

In addition, the small screen size of the D10 (4 by 5 inches, as compared to the 340's 9 inches square) seems a greater reduction than the spot size, again compared to the 340. This is a subjective judgement and depends critically on the D10's focus adjustment, but I feel confident in asserting that the D10 has not much better resolution, if any better at all. As screen and spot size are reduced, vibration effects such as false floor jiggling and building vibration may necessitate more solid display/camera coupling, and ultimately an optical bench.

The manual for the 5103N family is Tektronix part number 070-1143-00, price \$8.25; for the D10 is 070-1132-00, \$5.50. Both may be ordered from Tektronix, 244 Second Ave., Waltham; inquiries may be addressed to Mr. DeRemer there.



These photographs illustrate the use of the grey scale display slave system. The first is vision file WIRE2, displayed with $Z = (1000 - \text{data})/8$, thus achieving 64 grey levels. The second is vision file RESS, with $Z = (730 - \text{data})/8$, achieving 59 grey levels. Both had center & magnify options selected, took about 1 minute to display on the 340, and were photographed at f/8 on 75 speed film. The third photograph is the same as the second, but with SNAP patched to simulate all the data having a value of 430; thus it should be a uniform grey. It illustrates deflection hum and jitter, and a reflection from a small light leak in the dropcloth used to exclude ambient light.

EXPOSURE GUIDE FOR DEC 340:

$$f = \frac{k}{r_1 + r_2} \frac{A N P}{F} (1 + I)^{3/2}$$

- f - f-number indicated on lens.
- k - Empirically found to be about 1/125 (gives rise to density of about 2 in negative; i.e. almost overexposed).
- r₁ - Half-intensity radius of spot on DEC 340, varies somewhat with I. Use .5 mm unless you have good reason to suspect other value.
- r₂ - Half-intensity radius of blur in camera projected back onto display surface - varies with lens and film used.
- r₃ - Spacing of points in image you are displaying. Use ∞ if all the points can be resolved in the image
- s - Scale send to DEC 340, 0-3.
- A - ASA rating of film.

For Polaroid B/W: 3000
For 35mm TRI-X: 300
For 16mm TRI-X: 200

N - Argument to .NDIS; i.e. number of times points are displayed.

P - Packing factor.

1 for resolved points

$\max \left(1, \frac{r_1 + r_2}{r_3} \right)$ for one-dimensional sets of points (vectors, increments, characters)

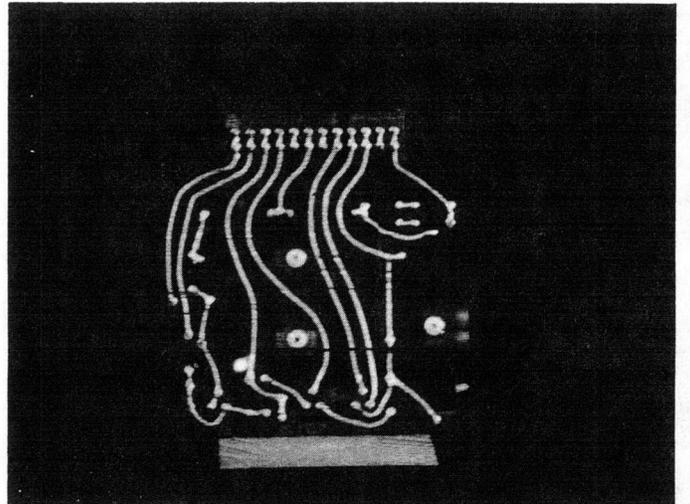
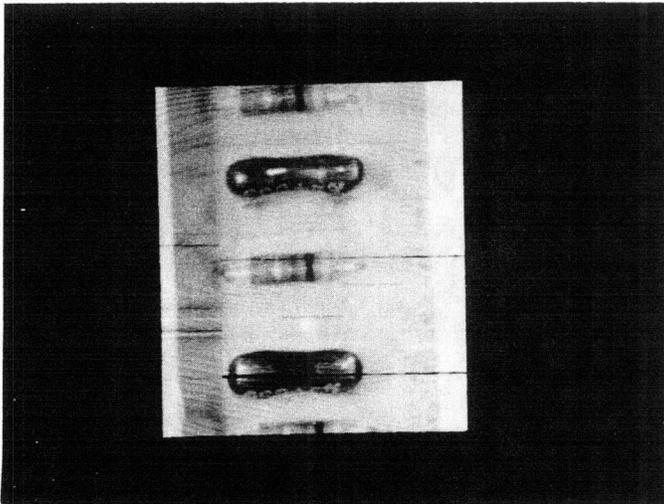
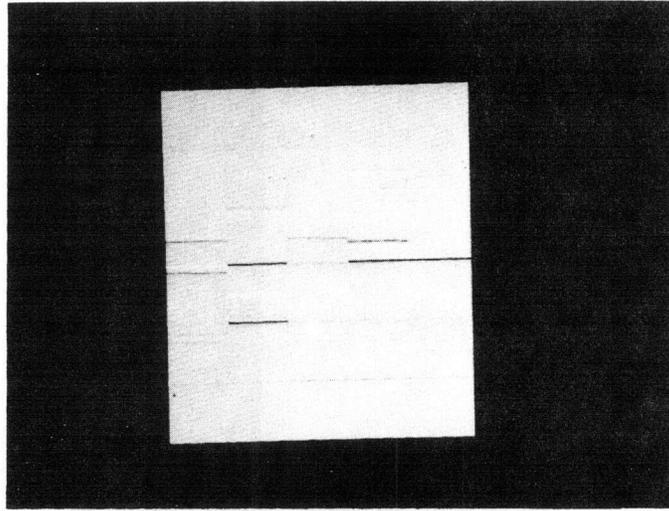
$\max \left(1, \frac{(r_1 + r_2)^2}{r_3} \right)$ for two-dimensional sets of points (rasters)

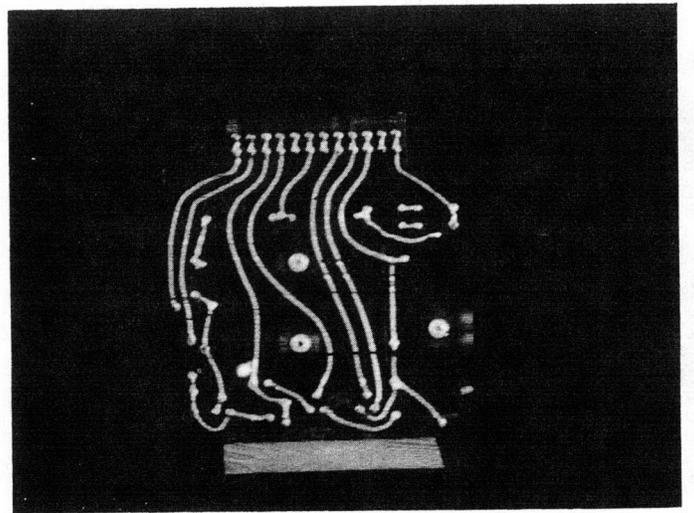
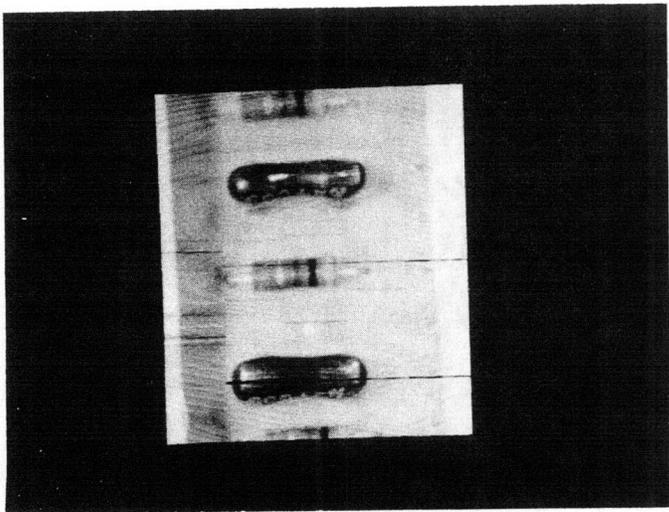
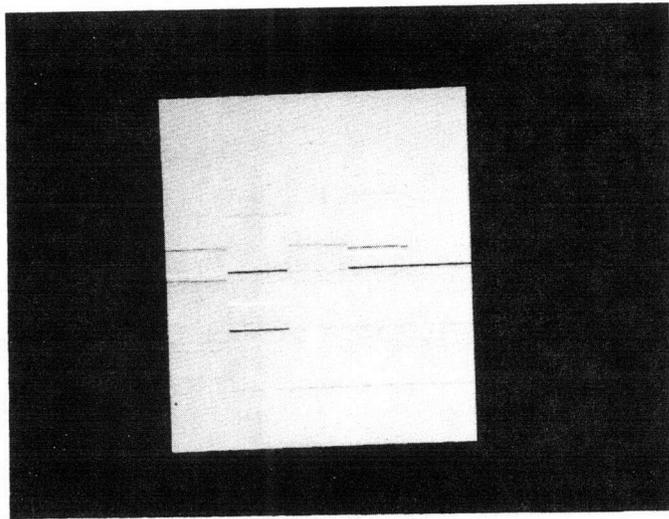
F - Filter factor. 1 for no filter, 2 for Wratten 15 (afterglow only), 8 for Wratten 47 (flash only).

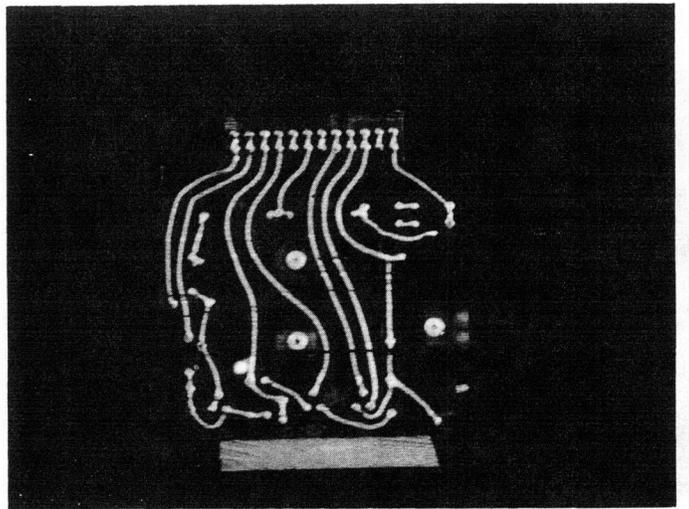
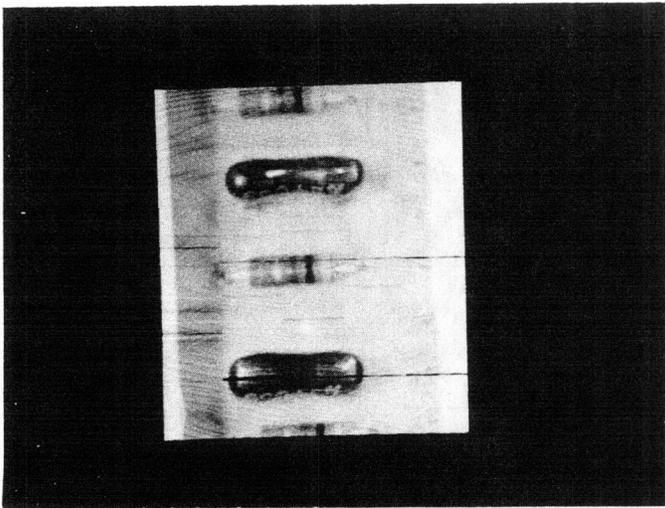
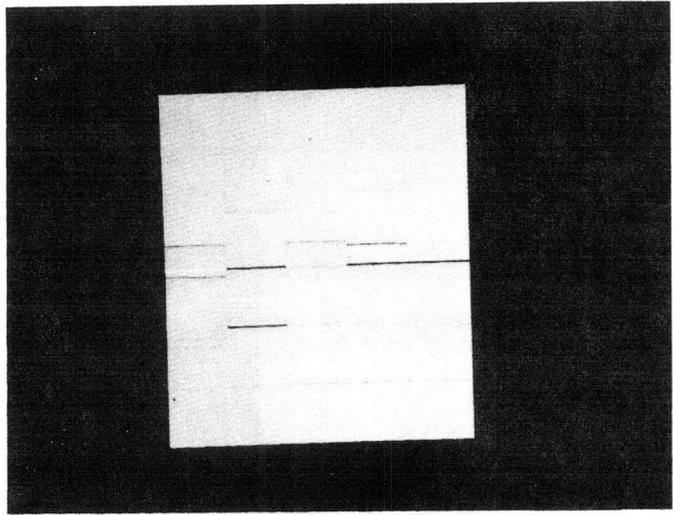
I - Intensity parameter send to scope. If varying intensities are to be recorded, use I = 5 - highlights will be slightly overexposed but the dark-areas will not be completely under-exposed. 0-7.

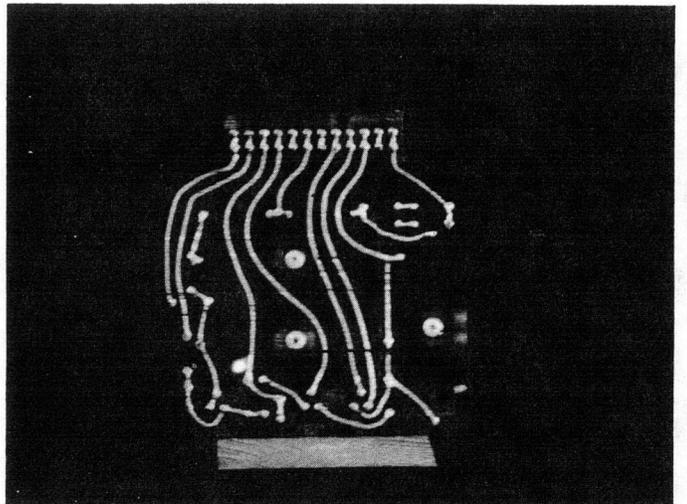
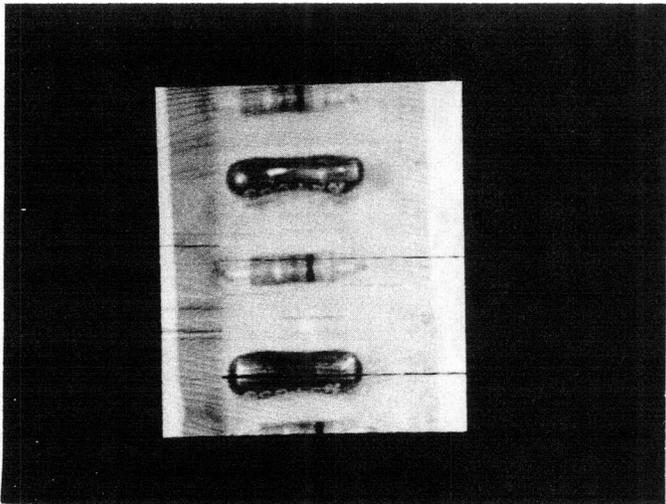
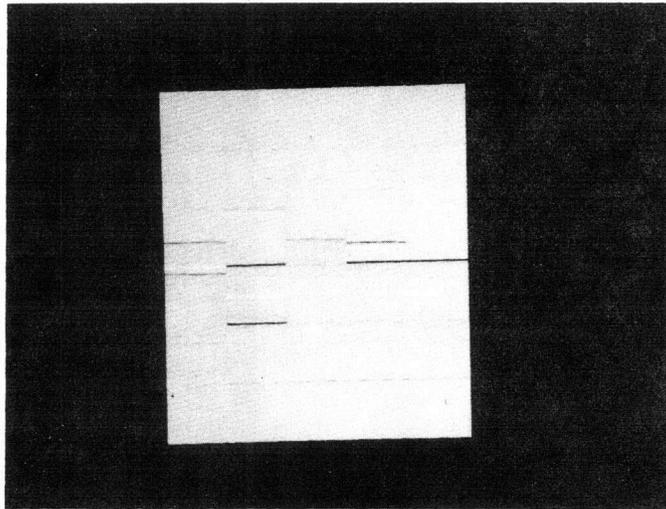
I (1 + I)^{3/2}

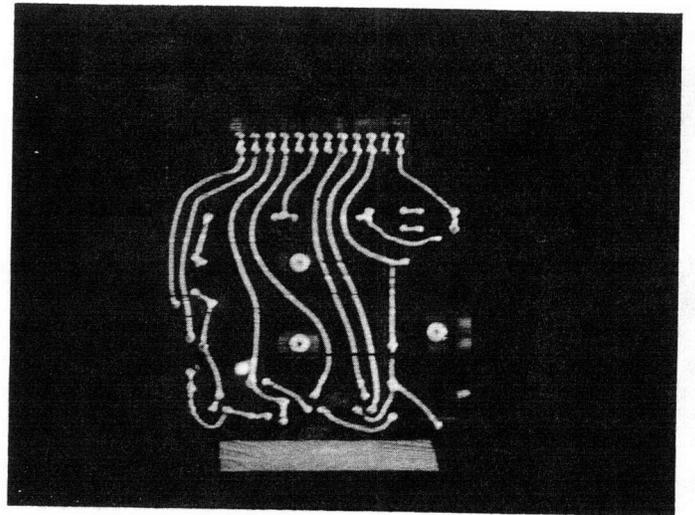
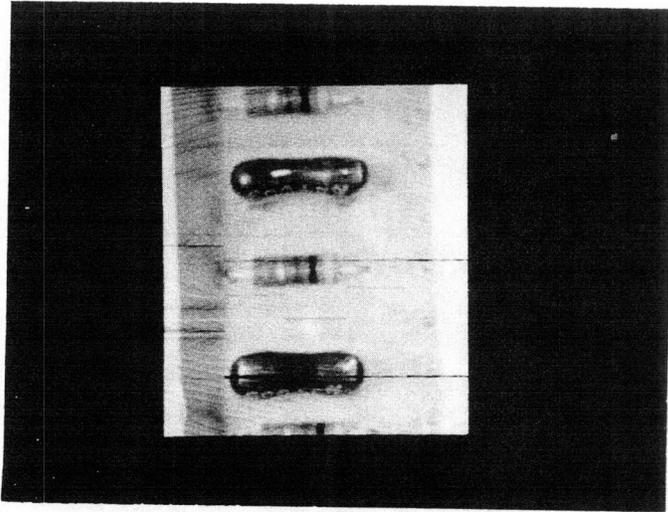
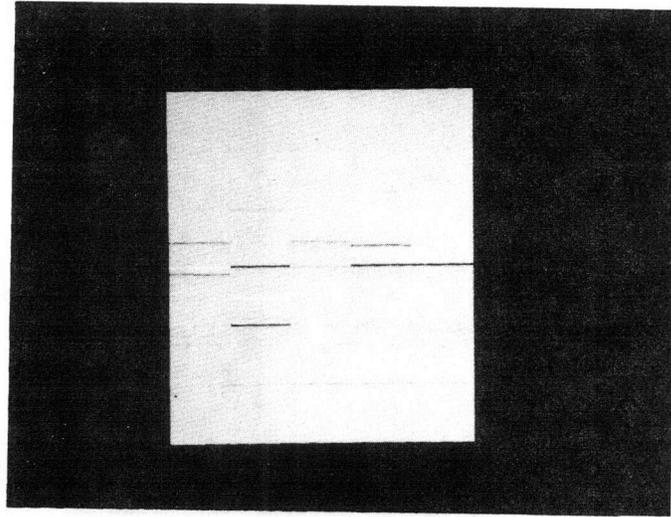
0 1
1 2.8
2 5
3 8
4 11
5 15
6 18.5
7 22.5

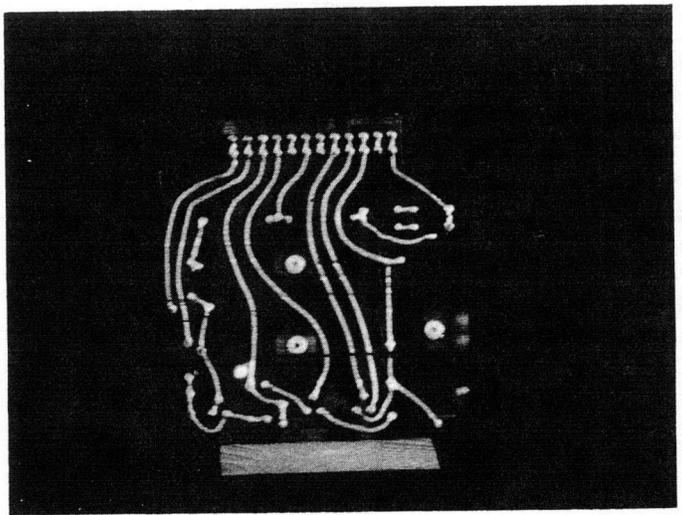
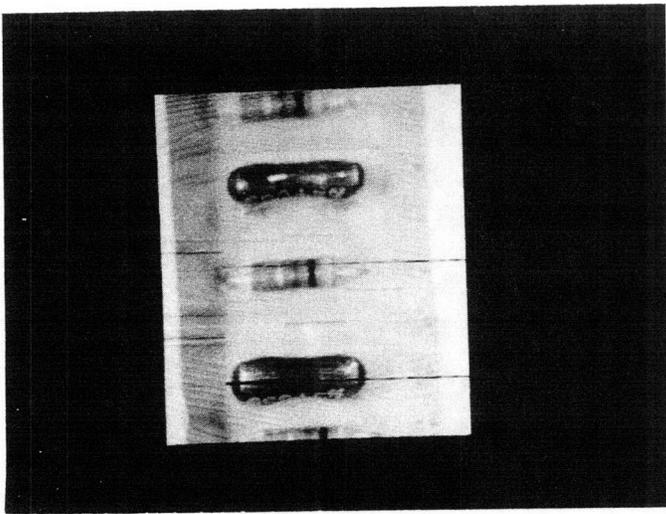
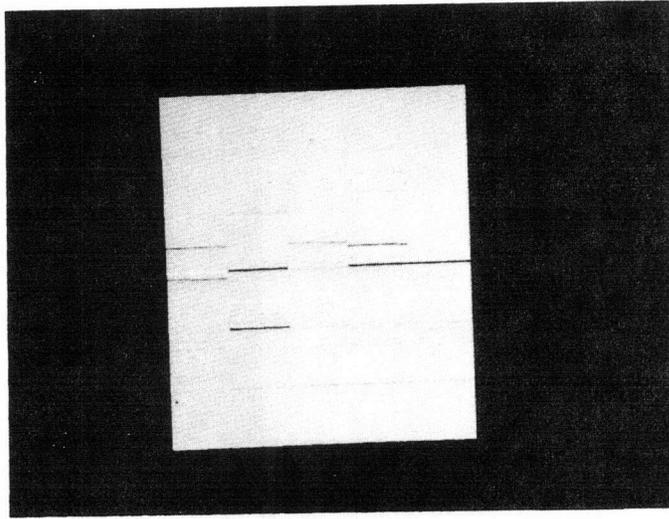


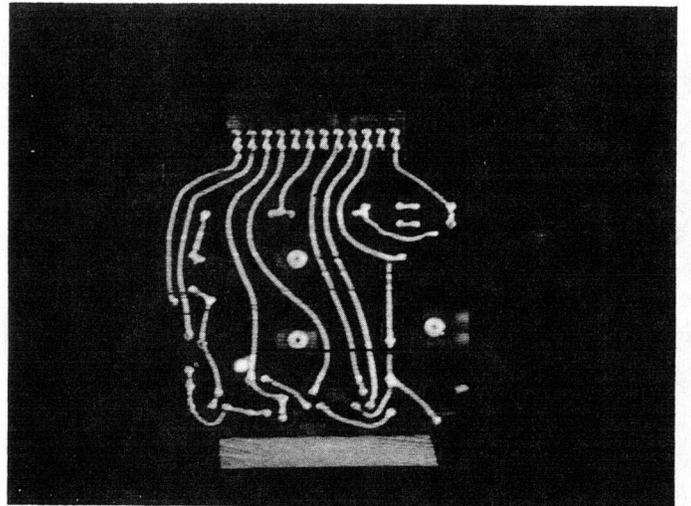
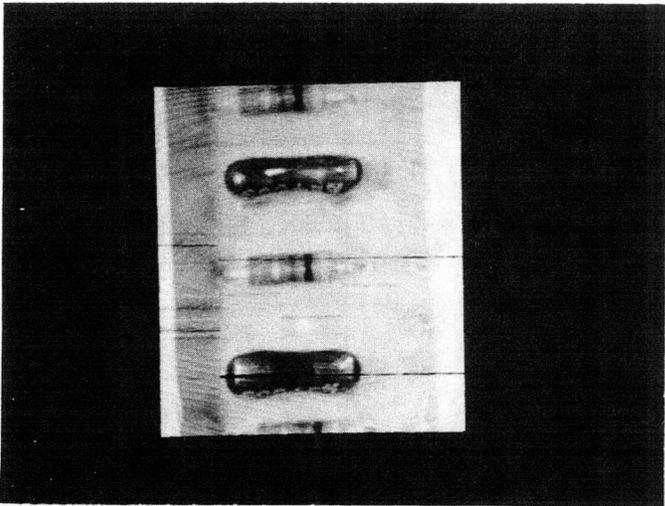
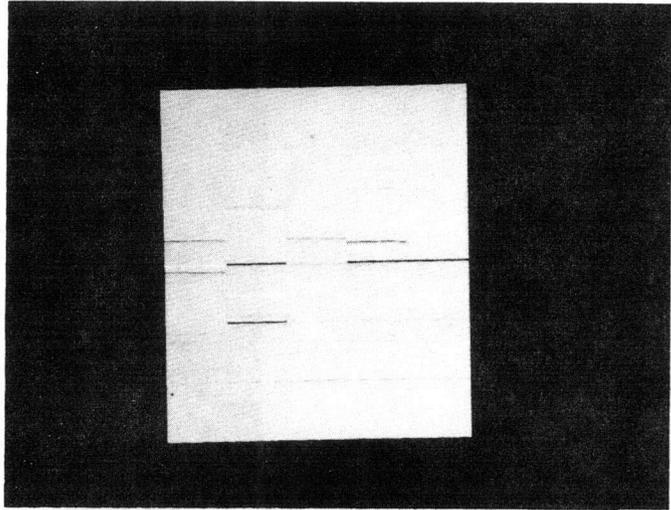


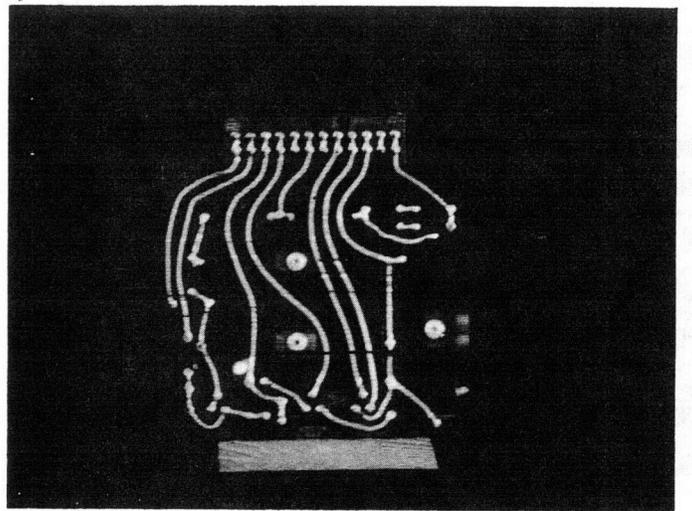
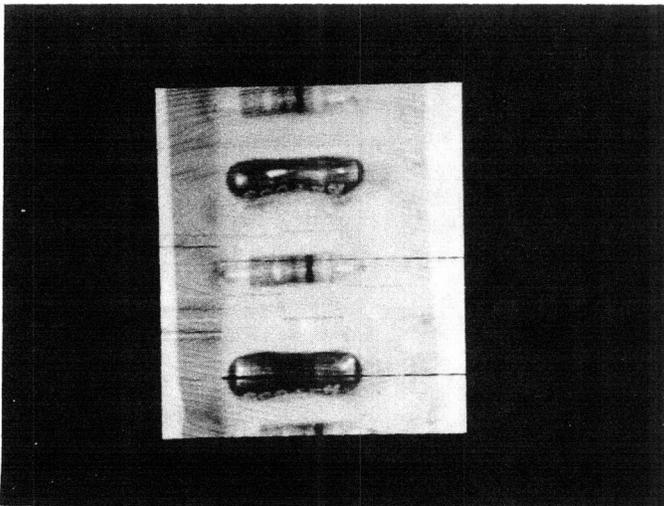
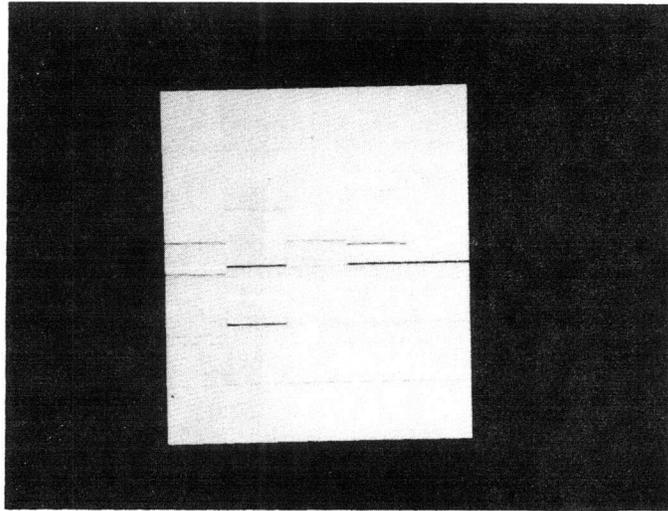


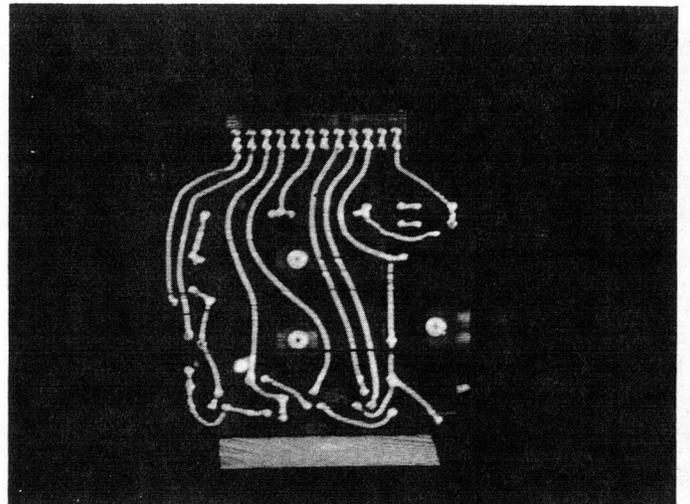
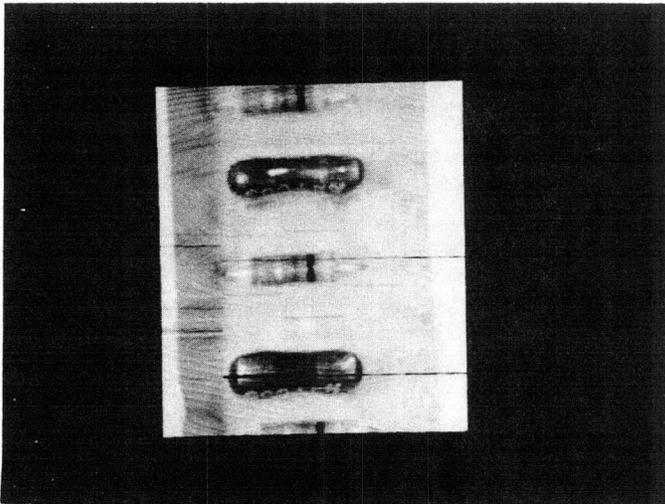
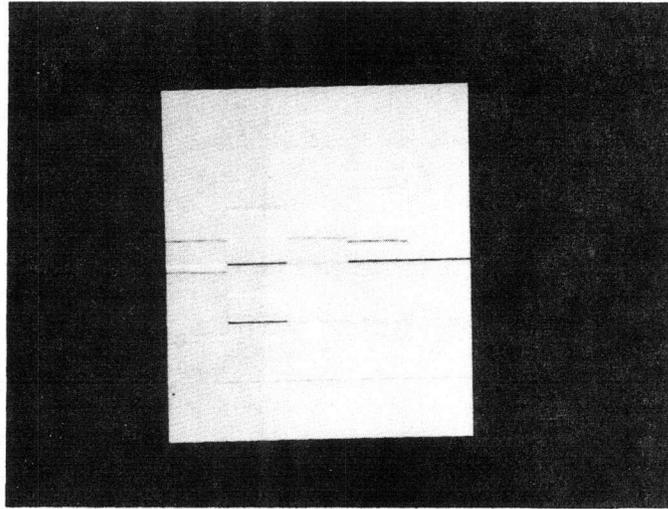


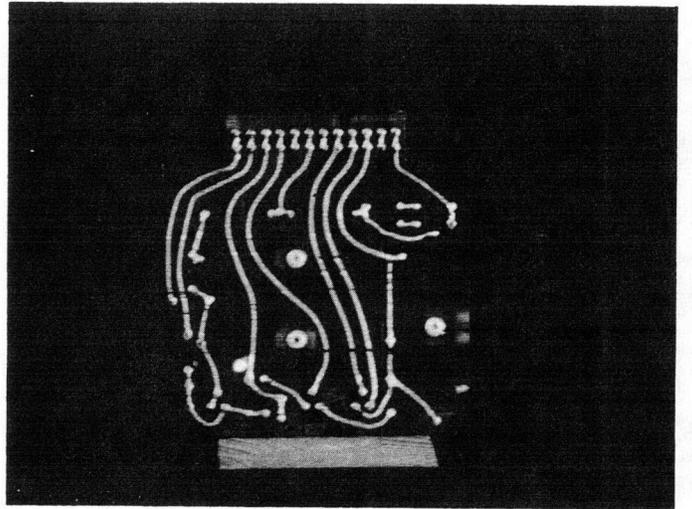
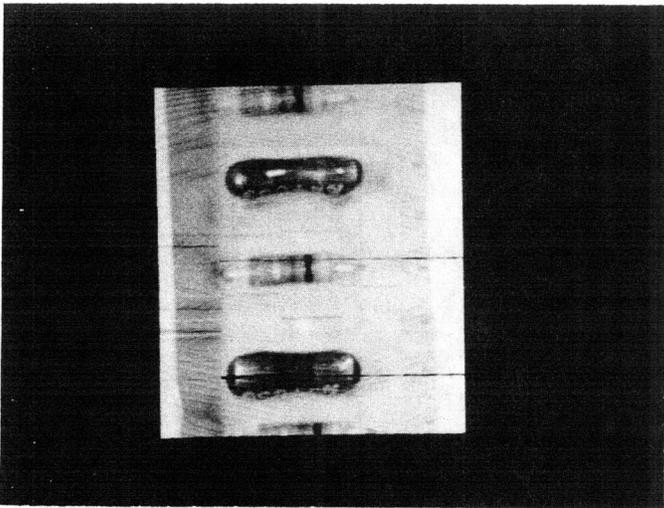
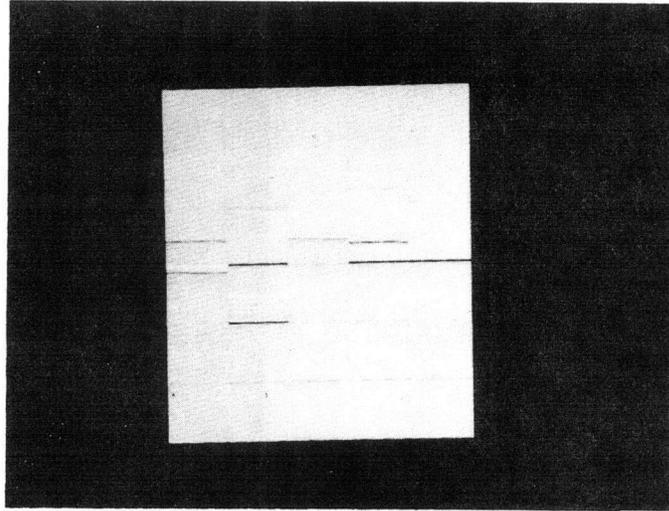


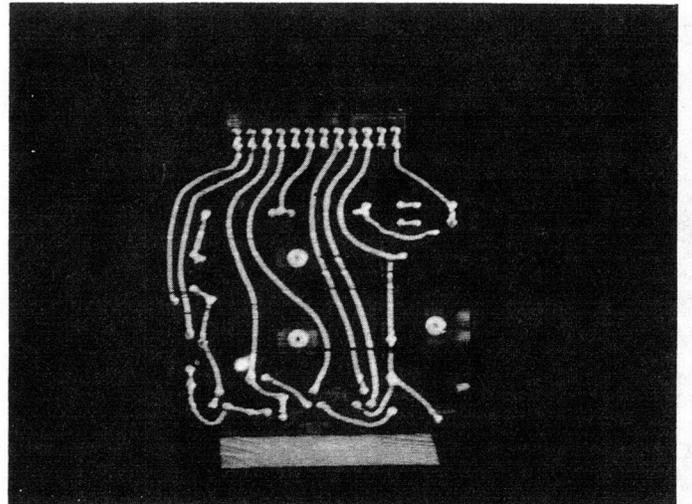
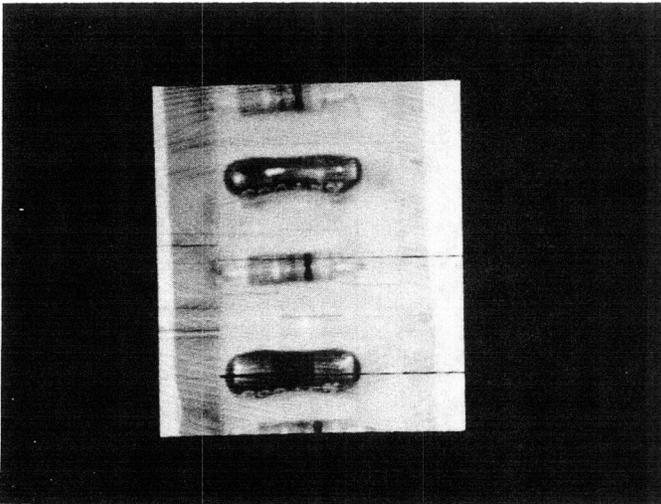
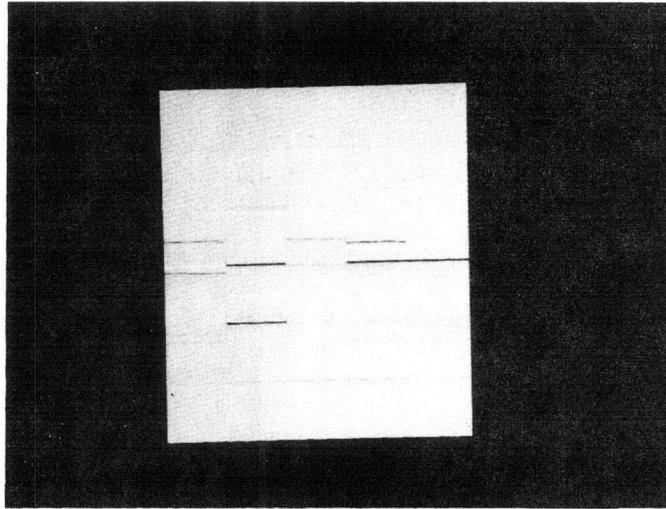


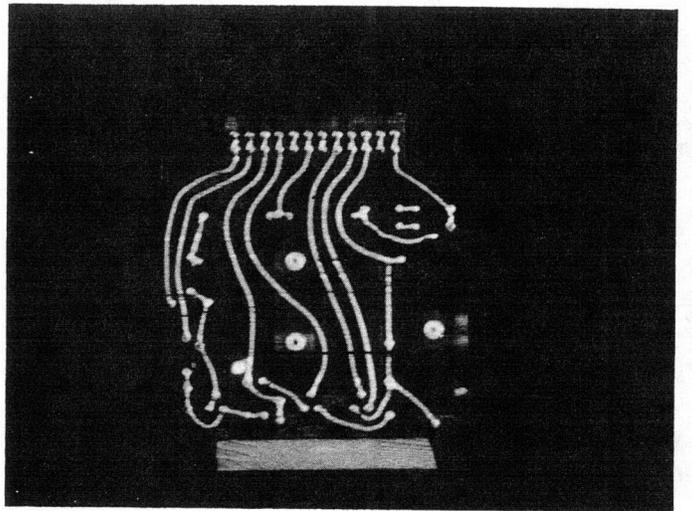
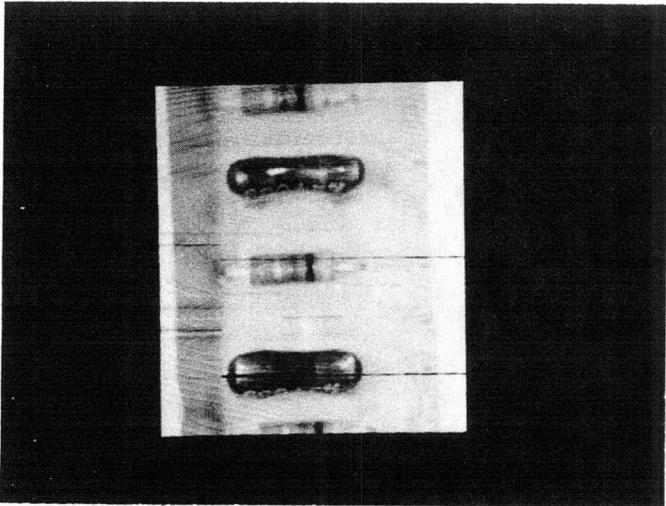
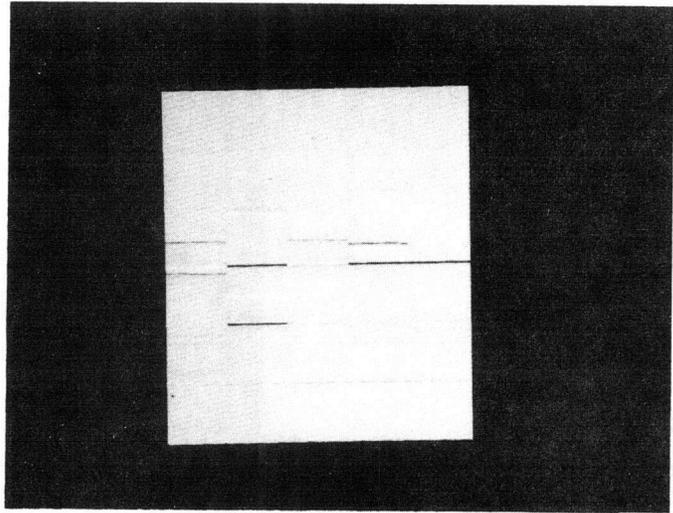


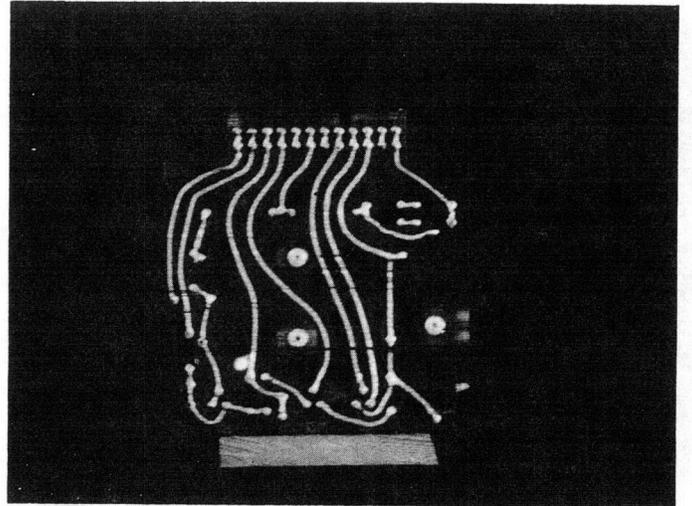
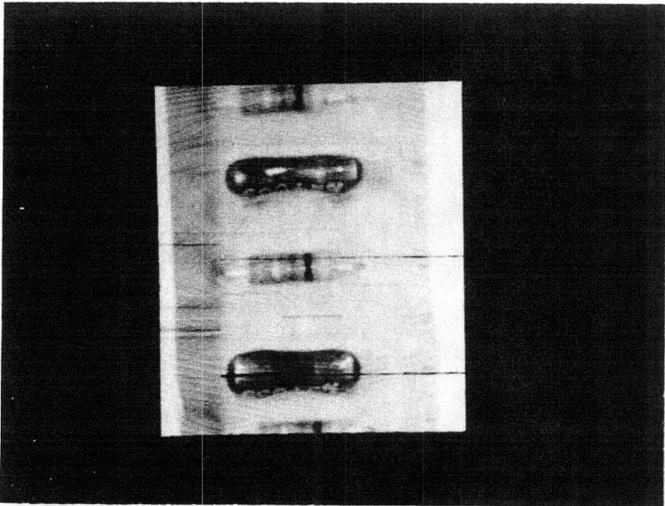
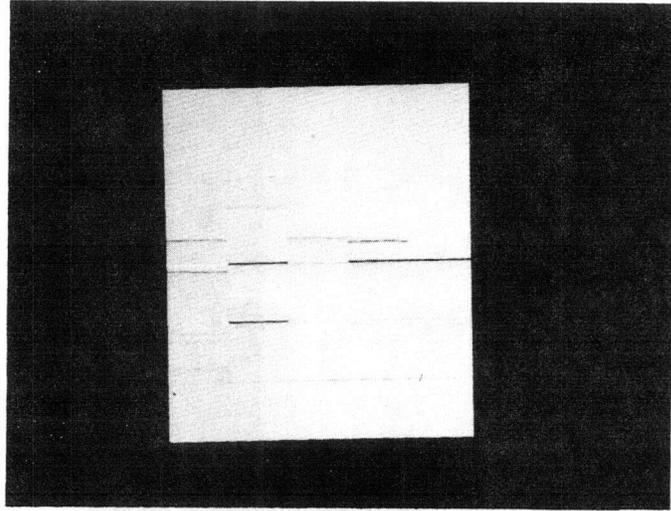


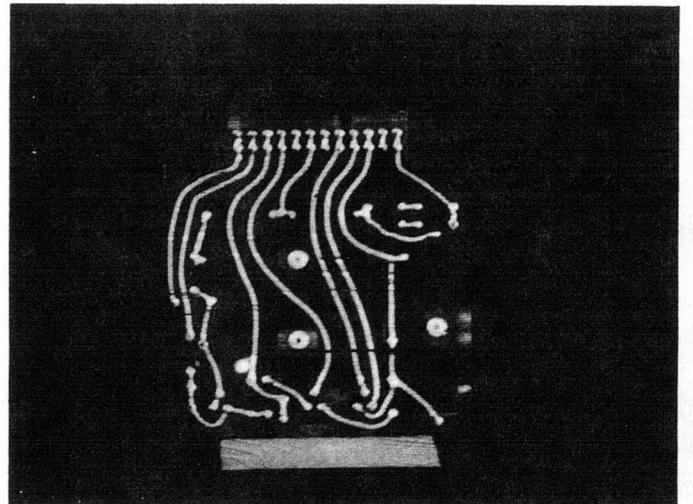
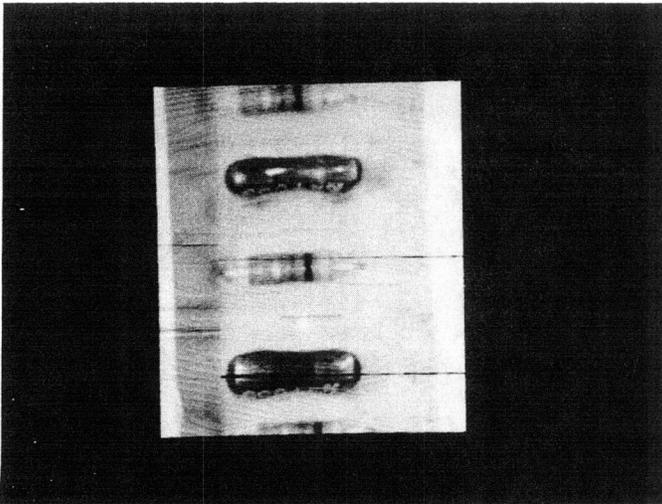
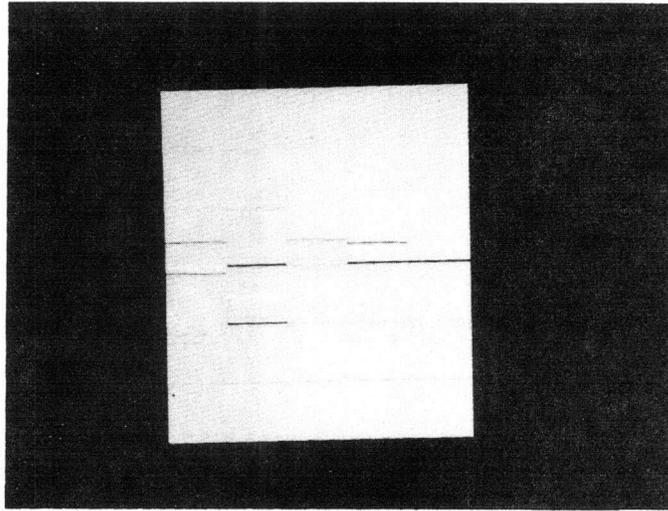


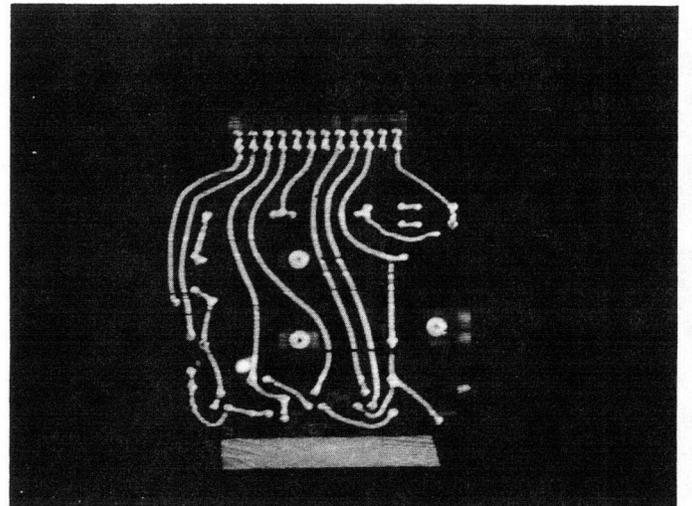
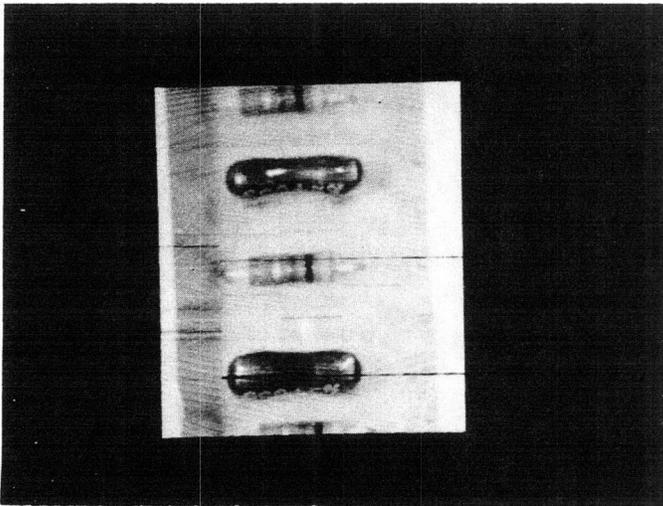
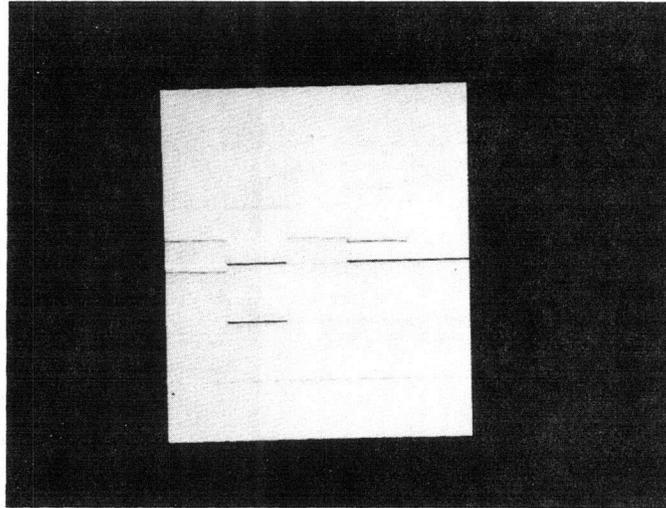


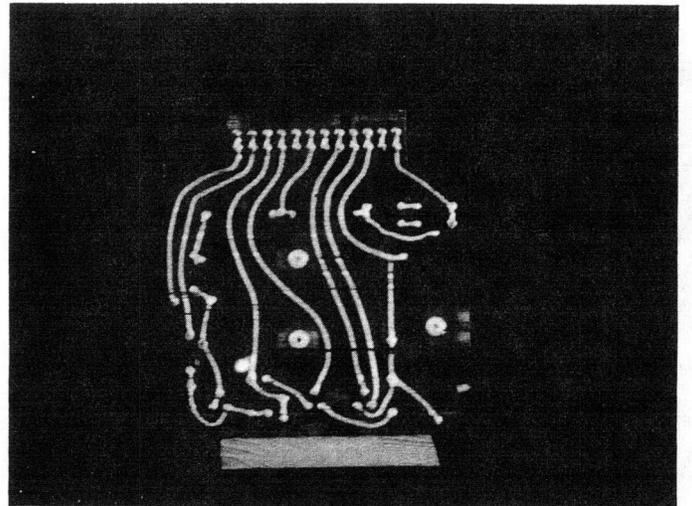
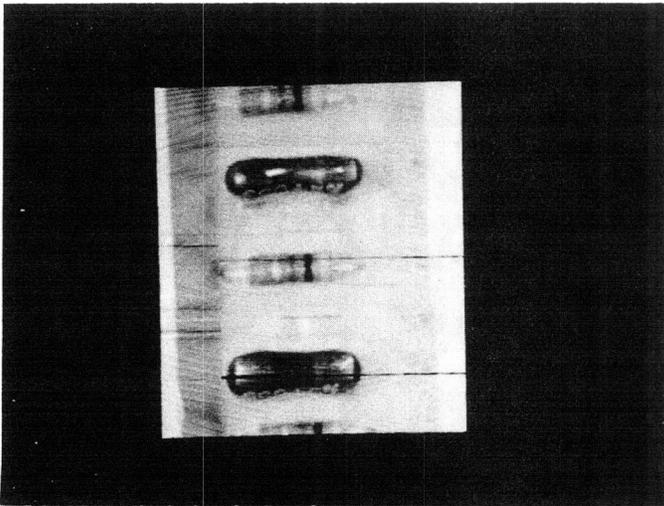
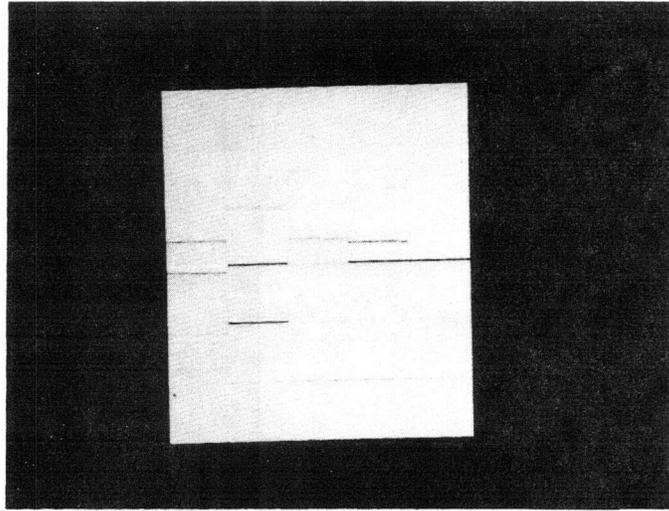


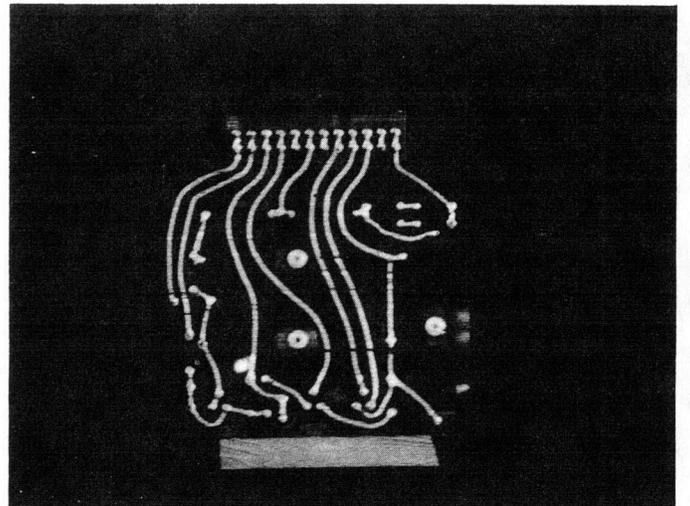
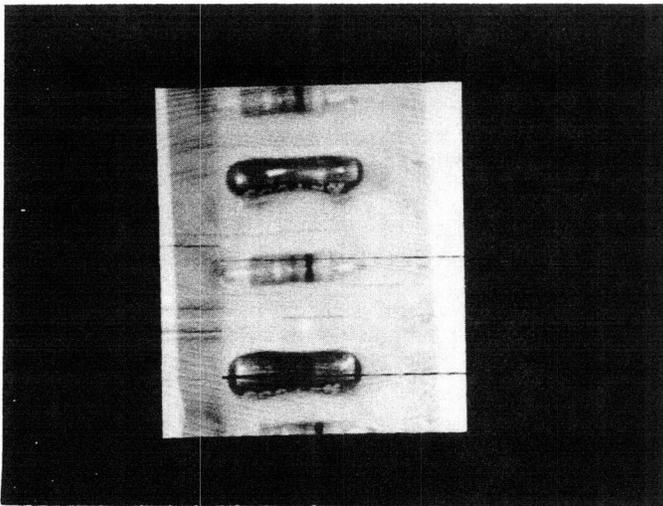
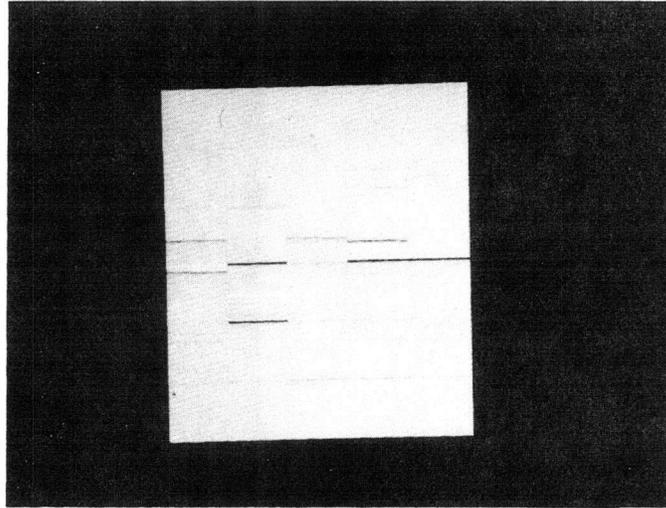


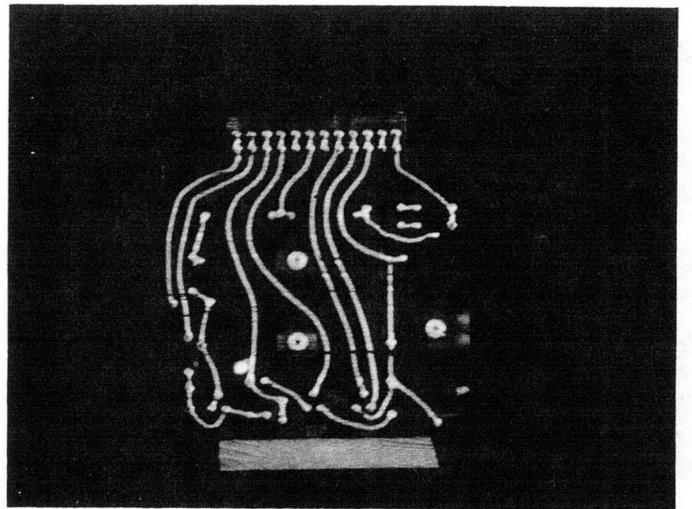
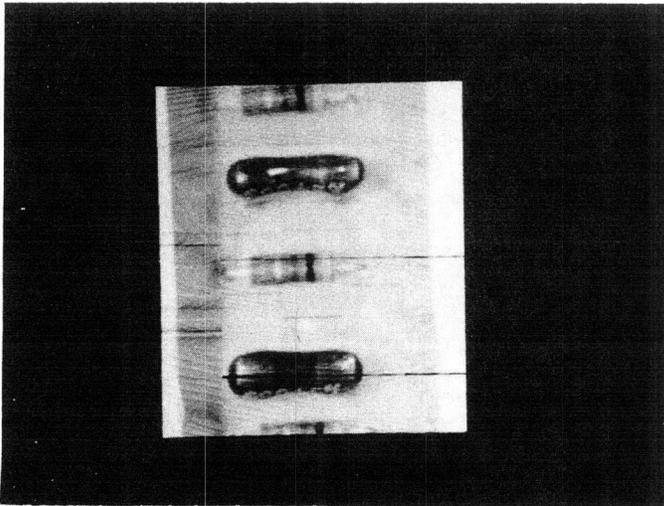
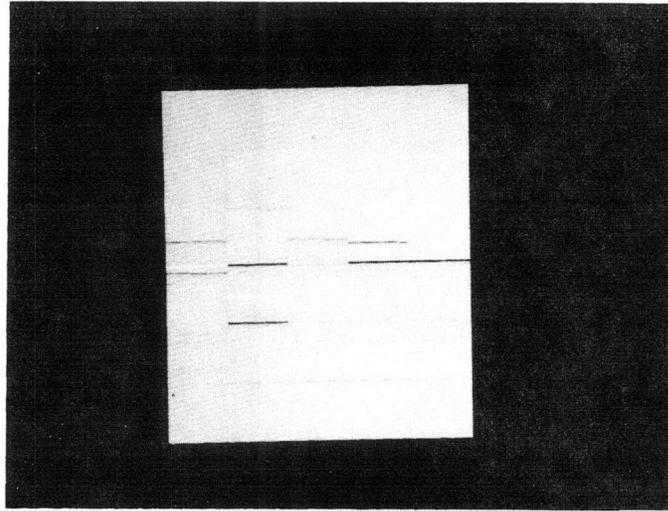


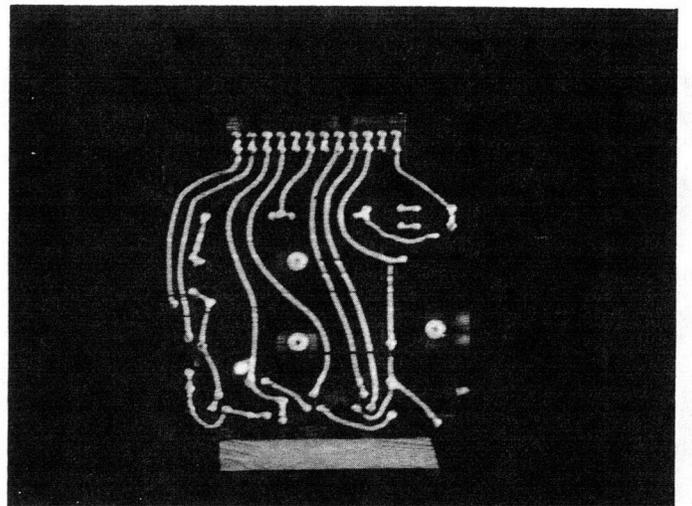
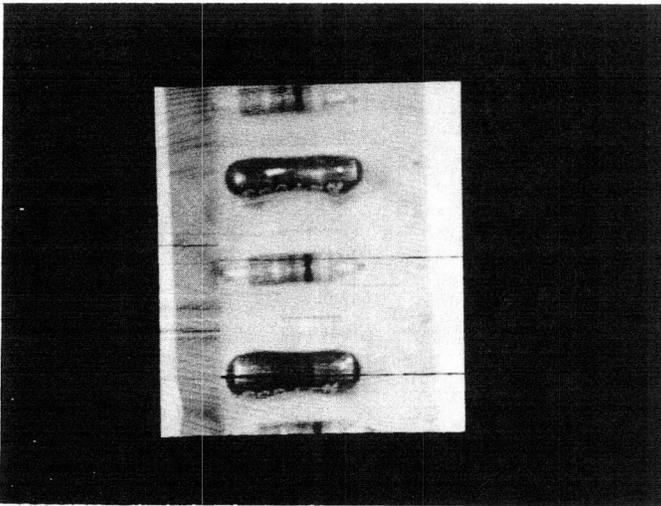
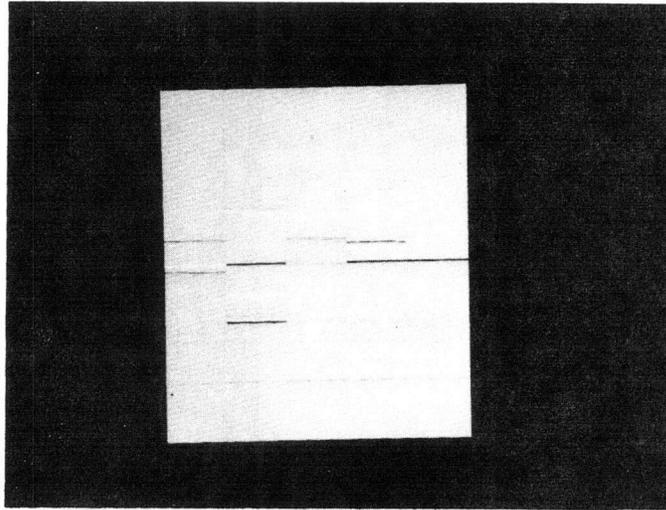


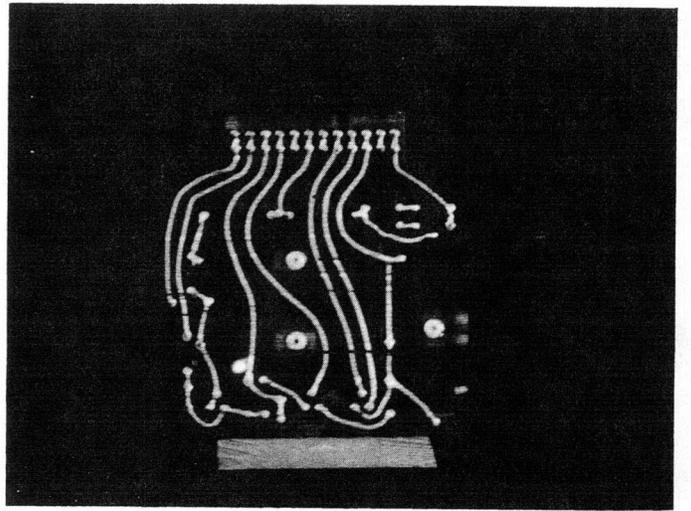
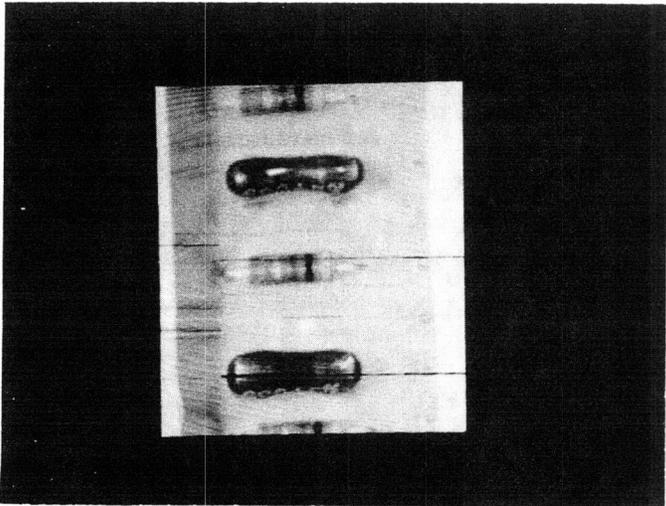
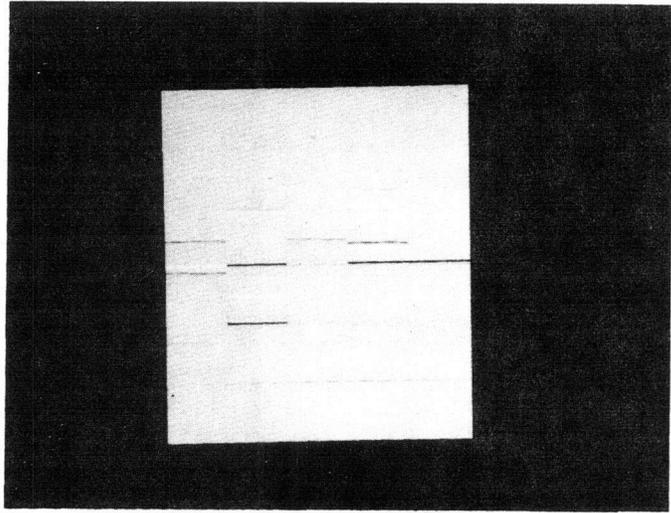


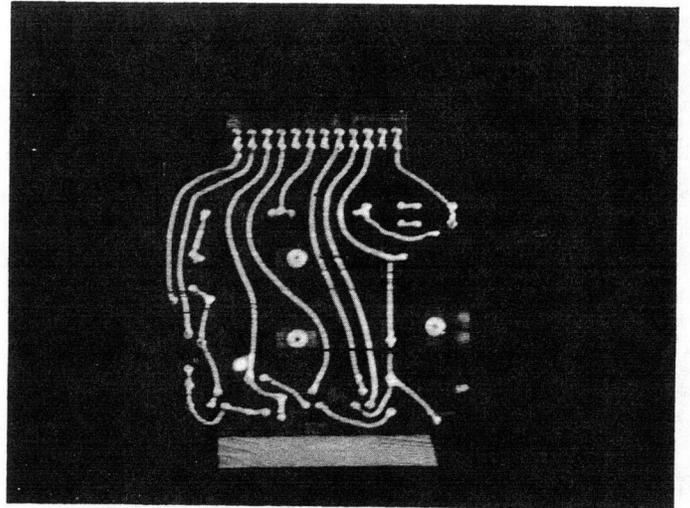
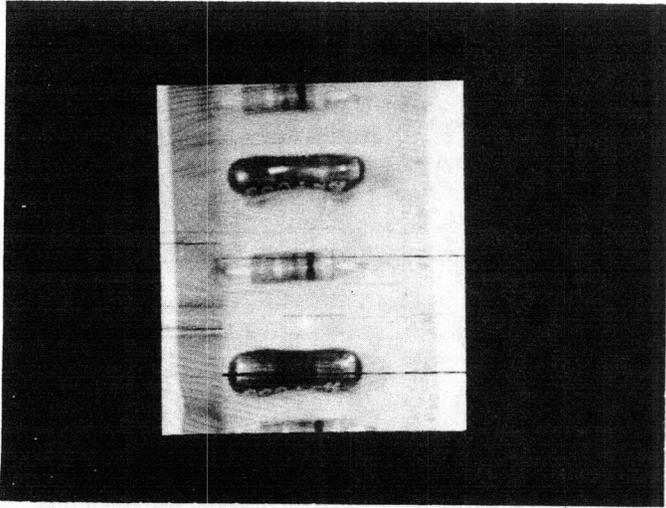
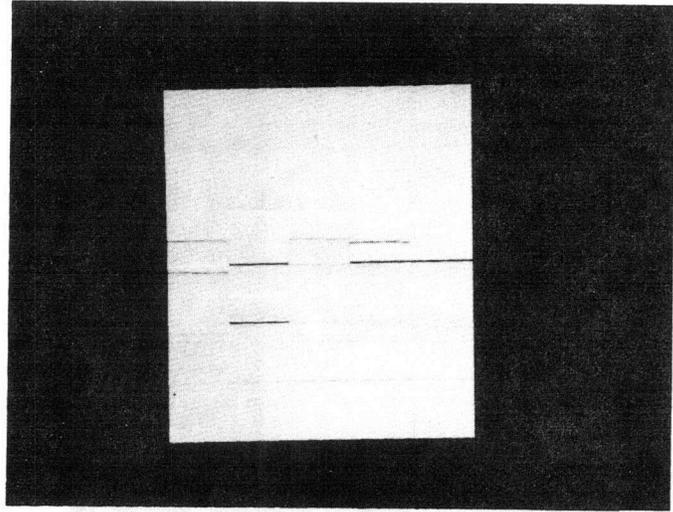


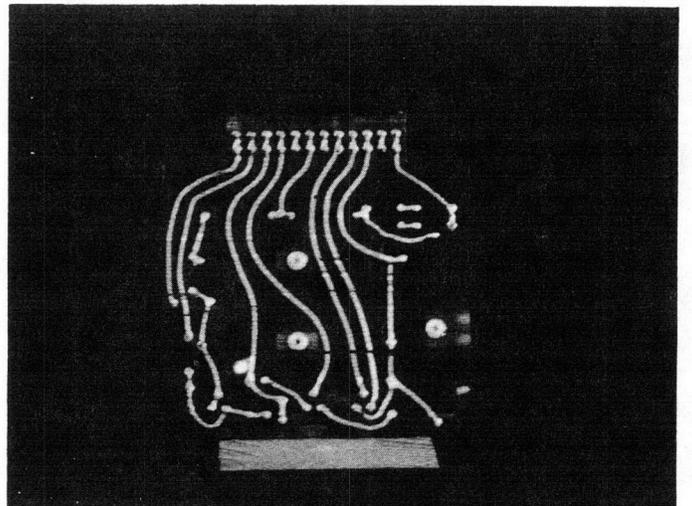
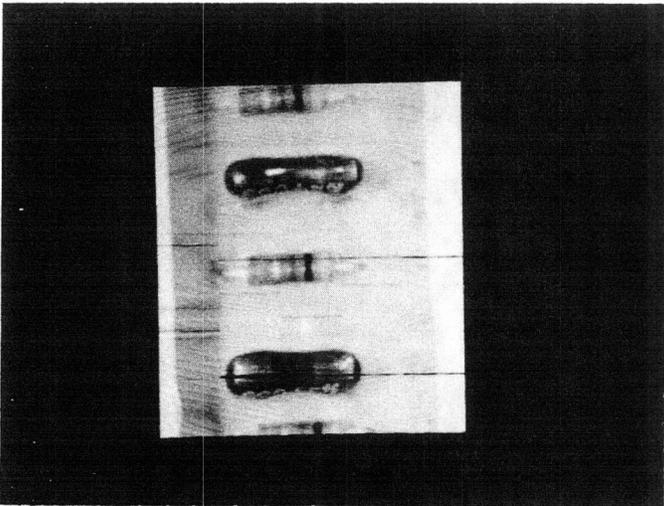
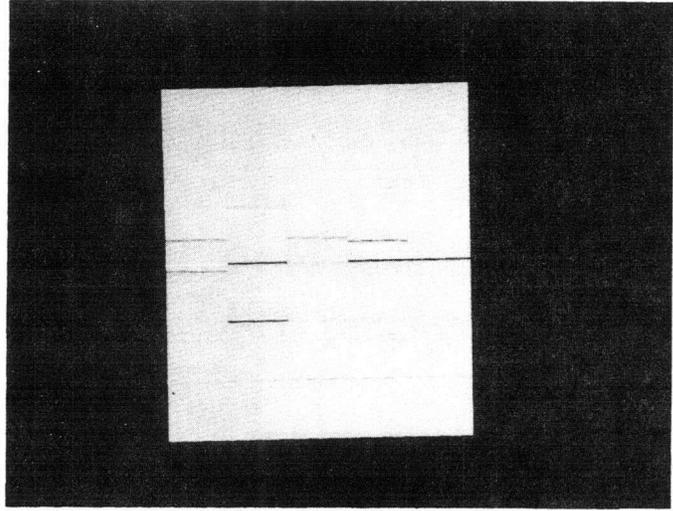


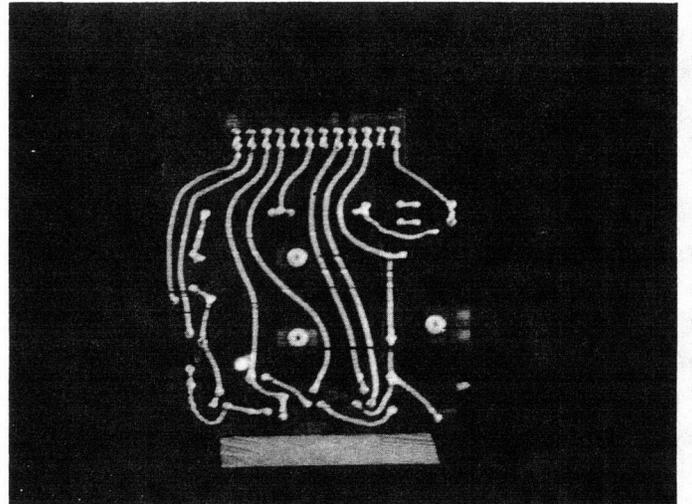
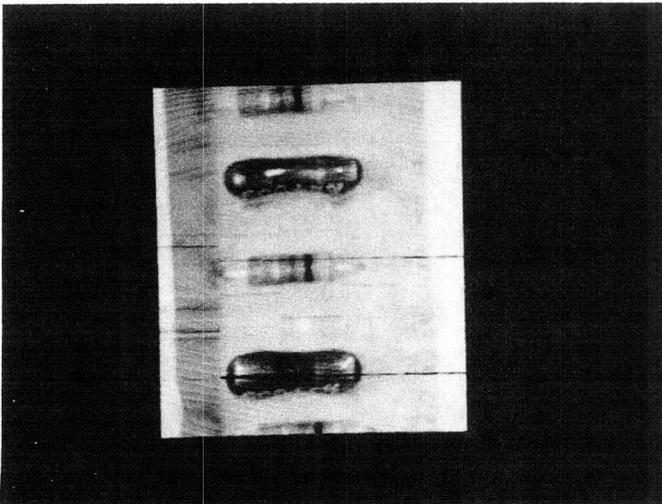
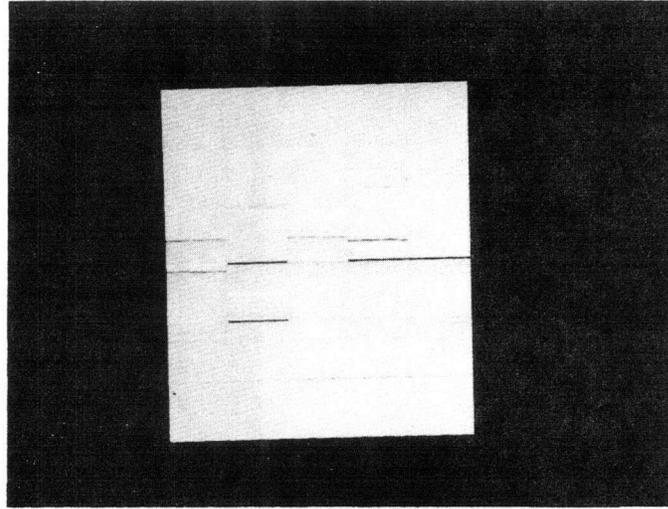


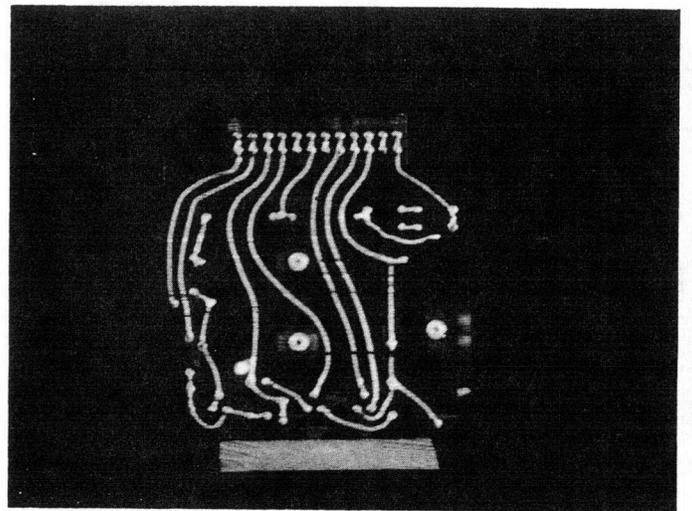
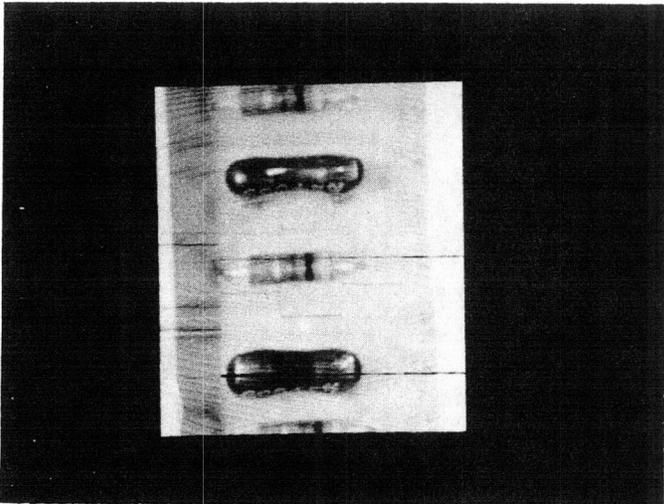
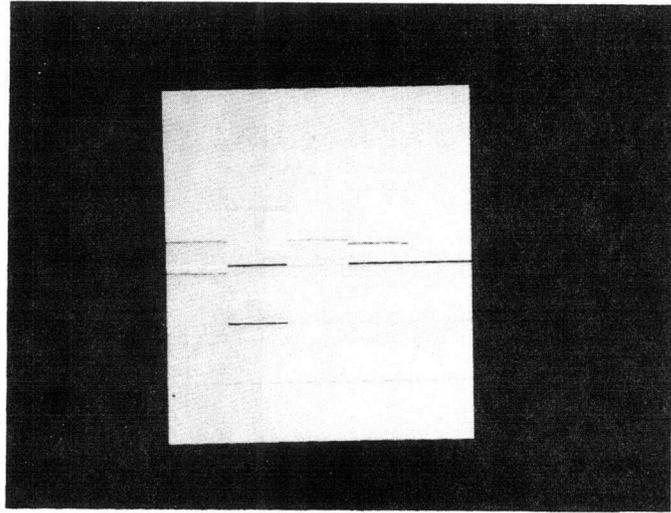


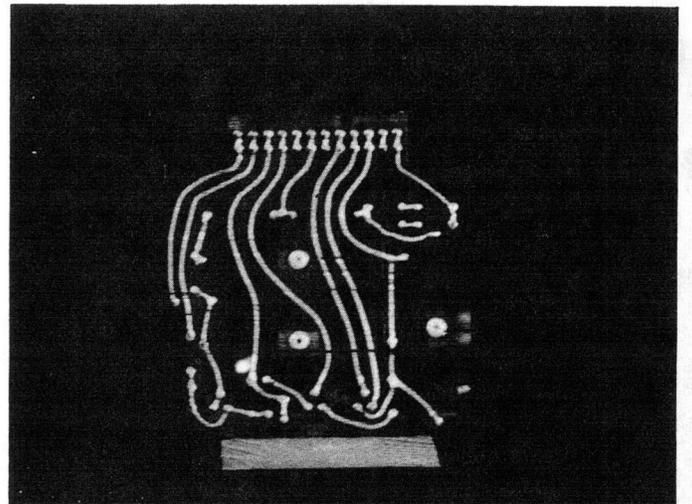
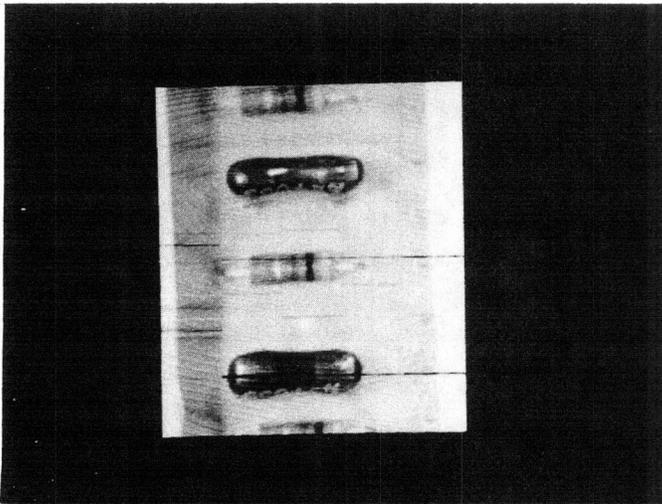
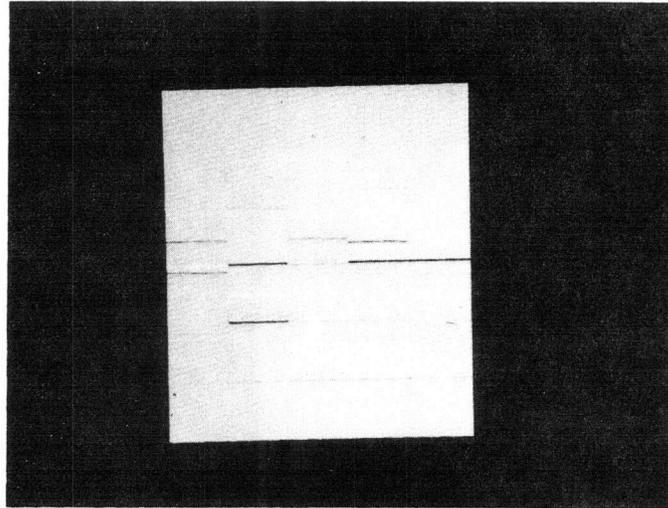


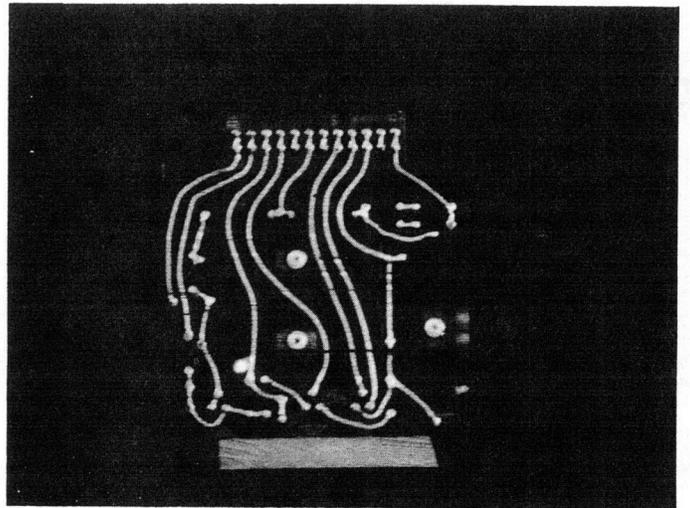
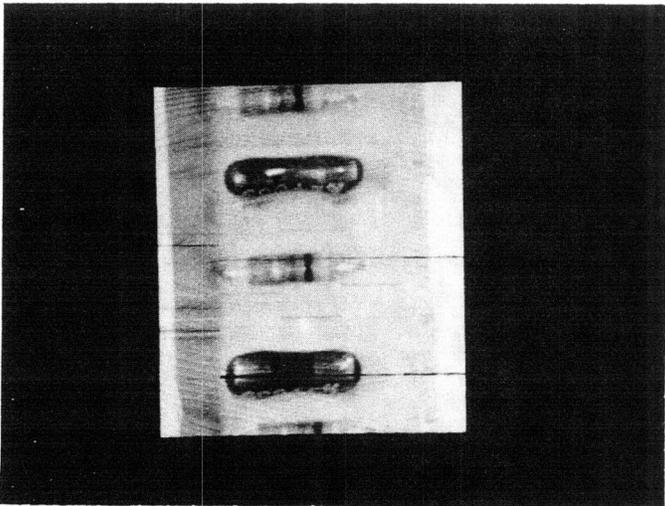
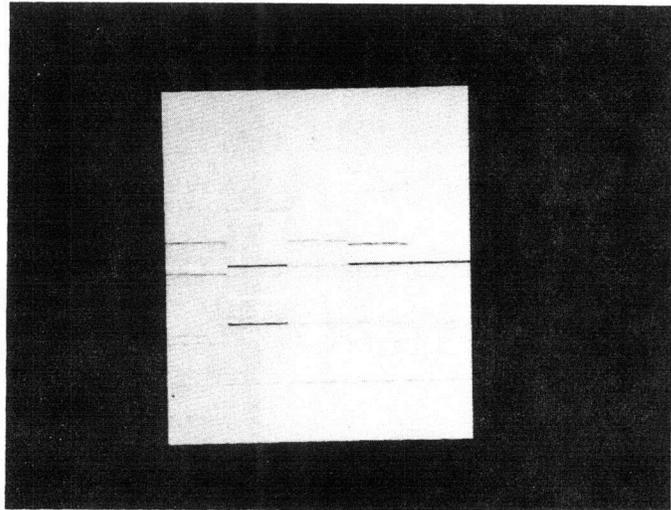


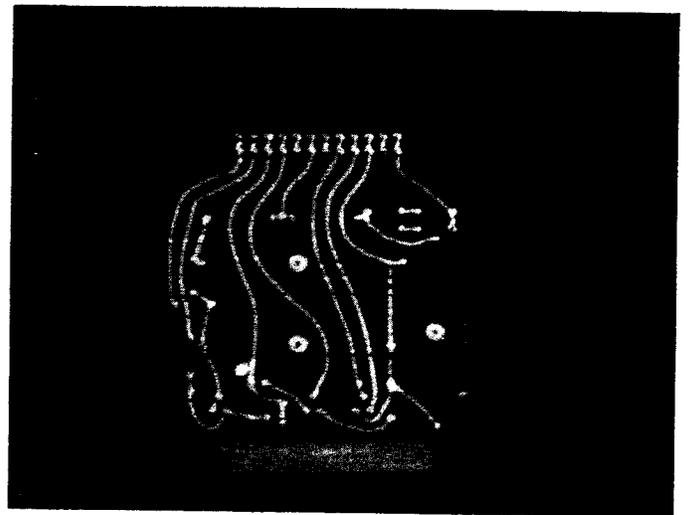
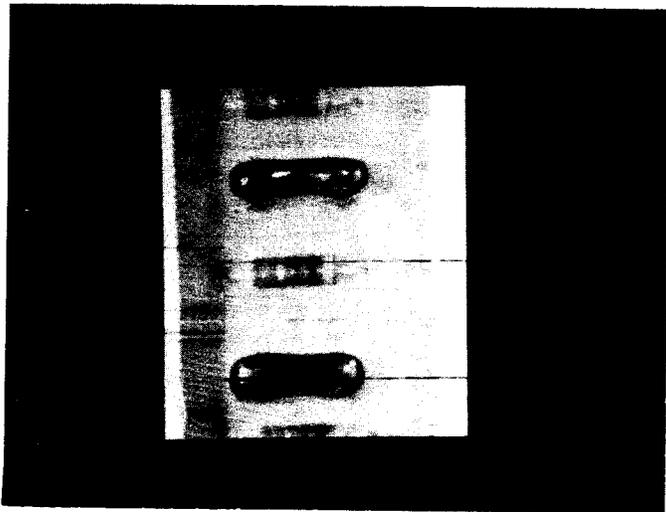
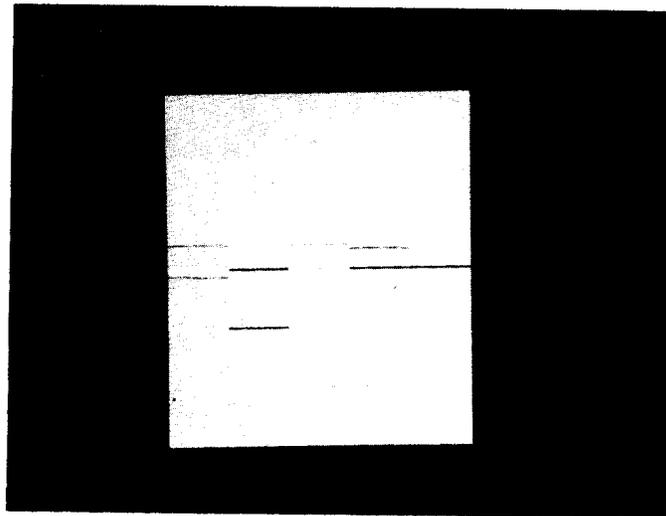


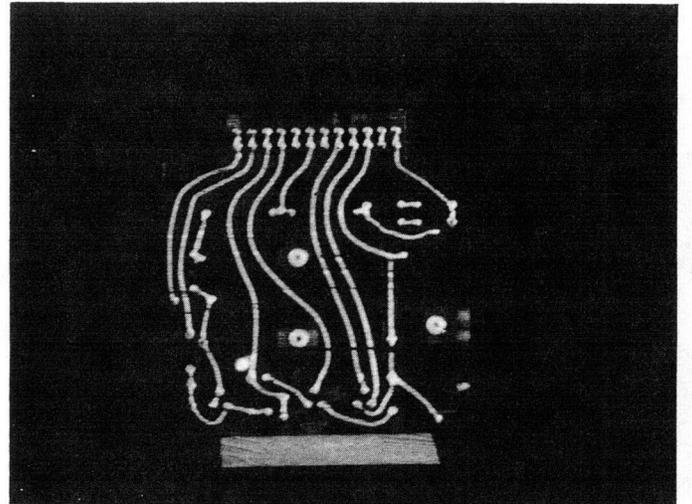
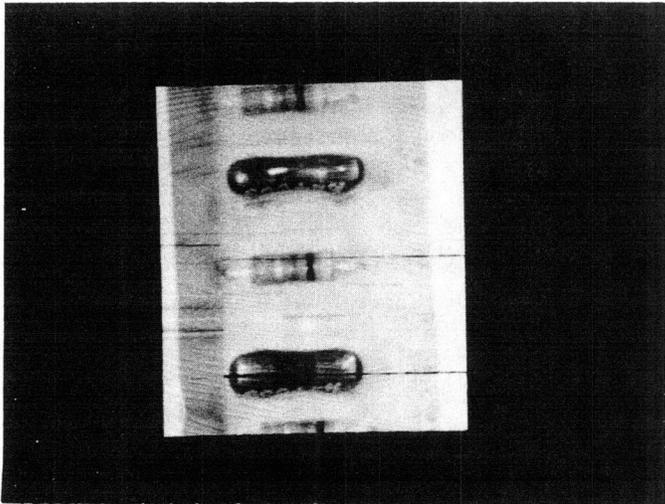
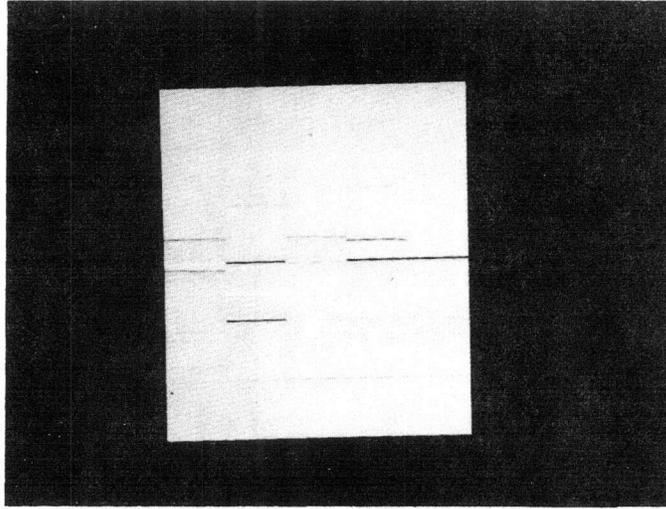




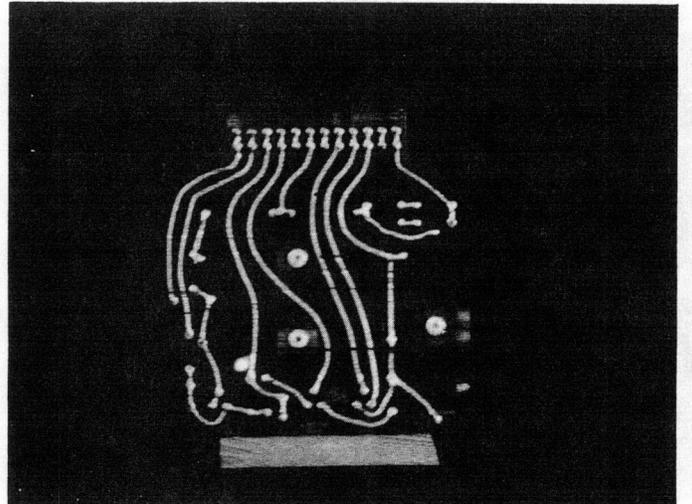
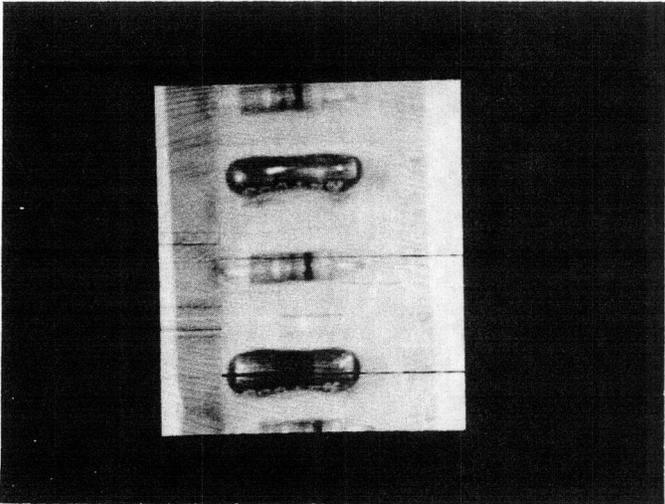
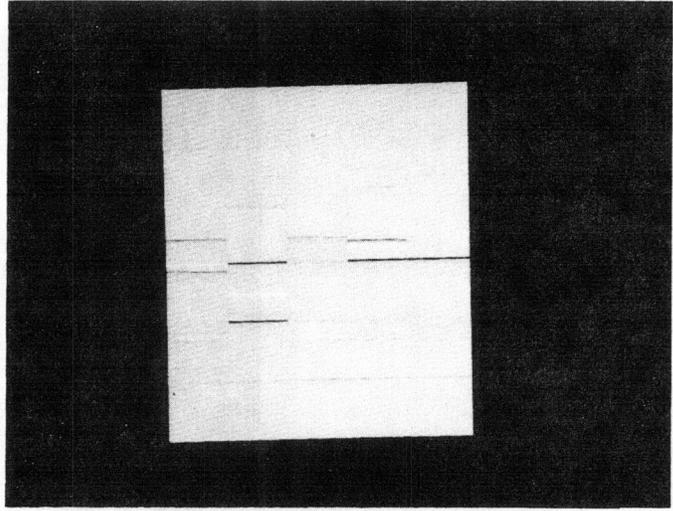


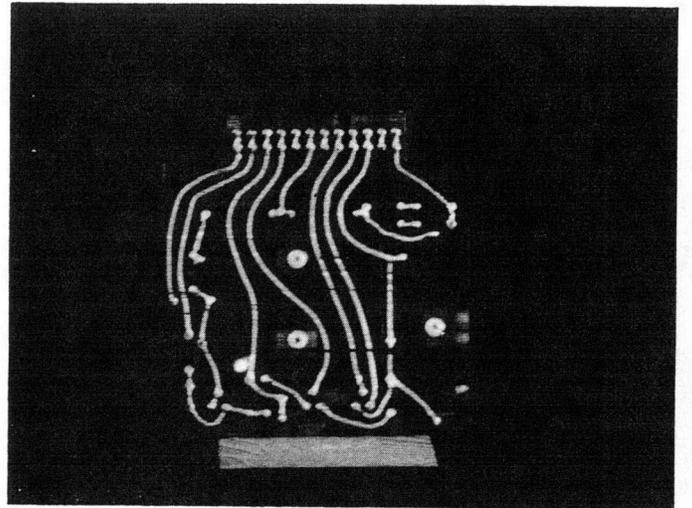
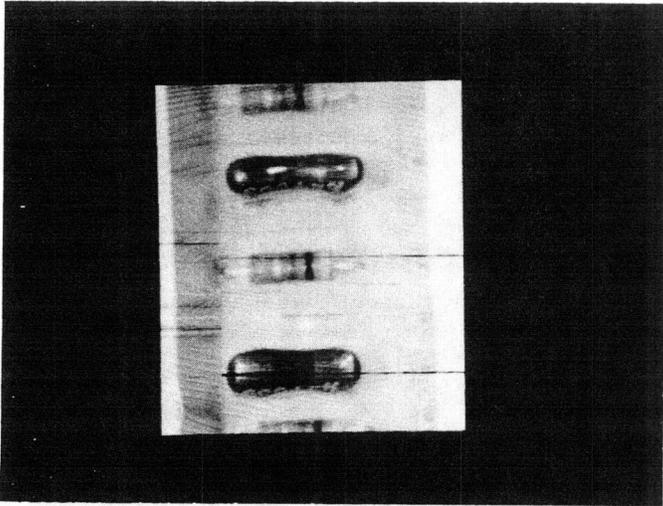
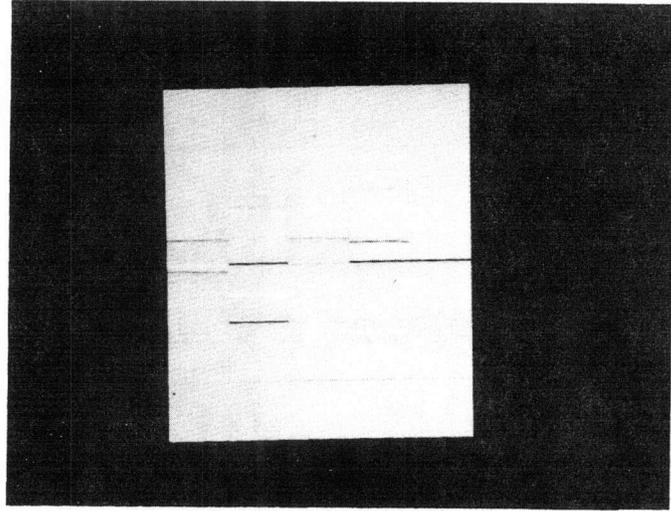




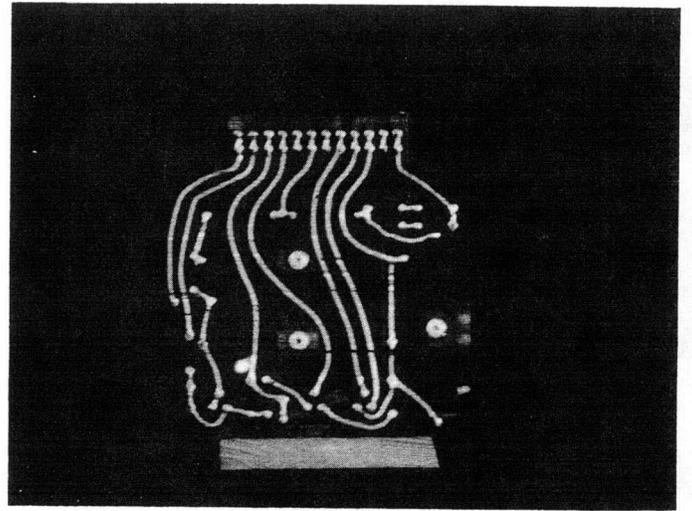
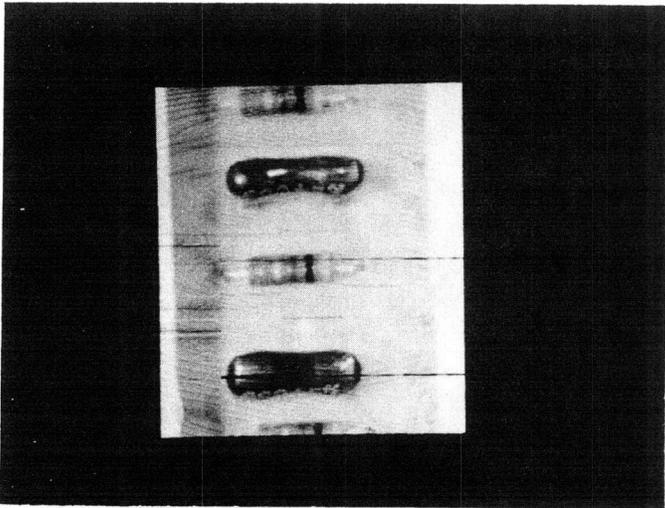
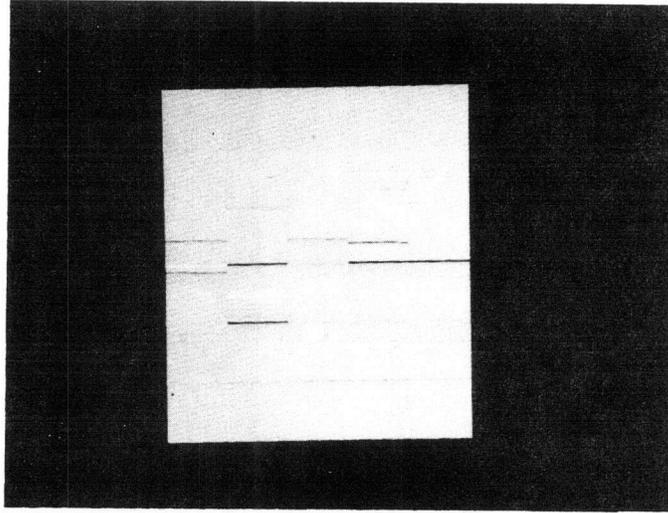


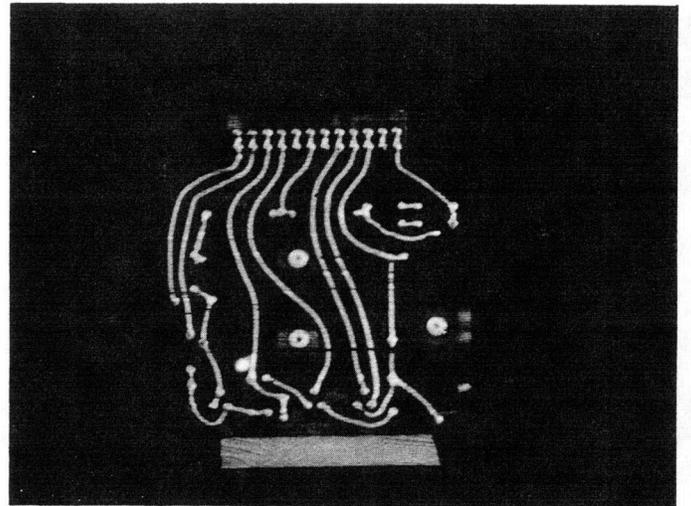
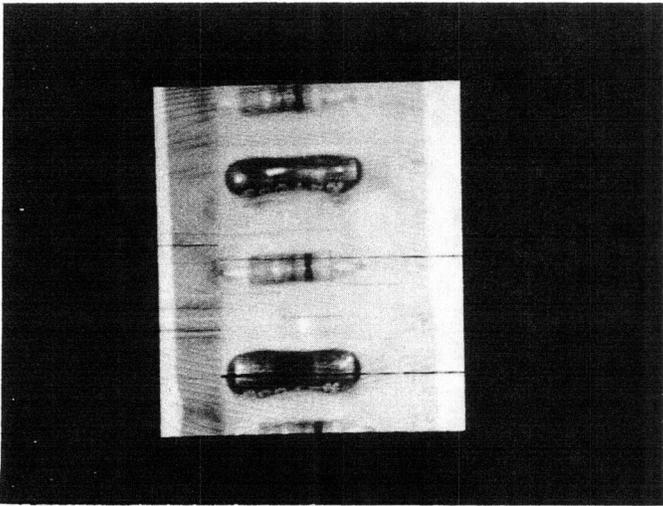
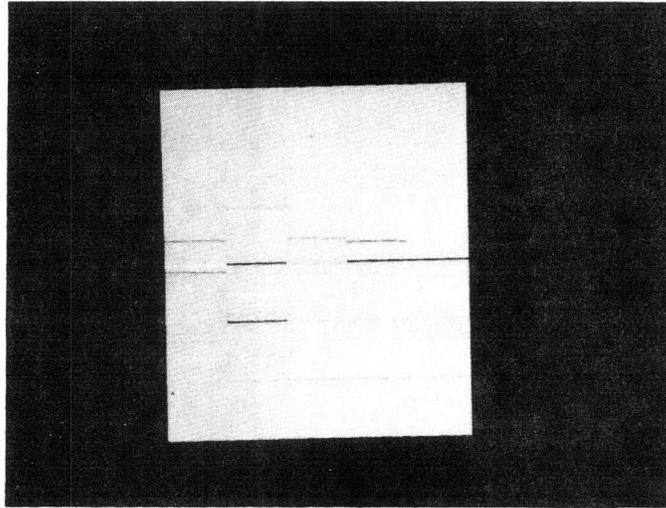
C₂

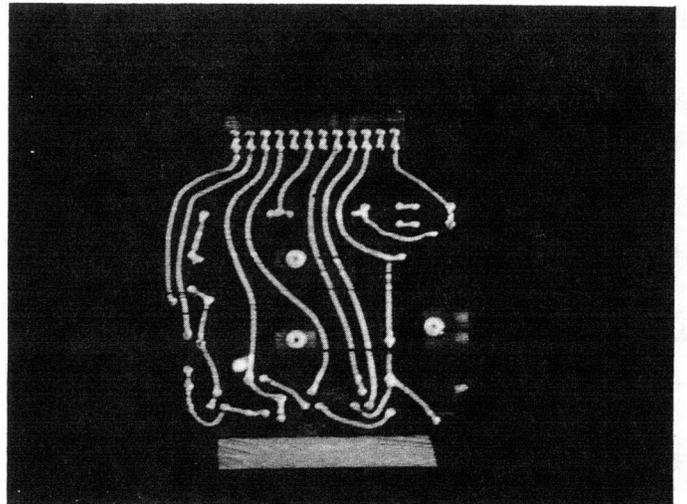
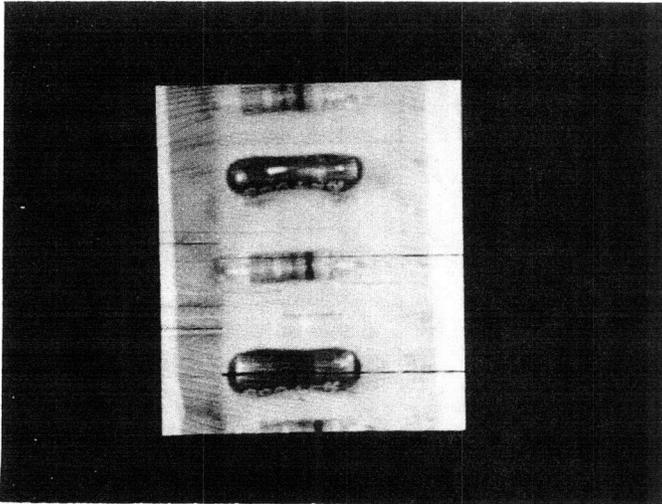
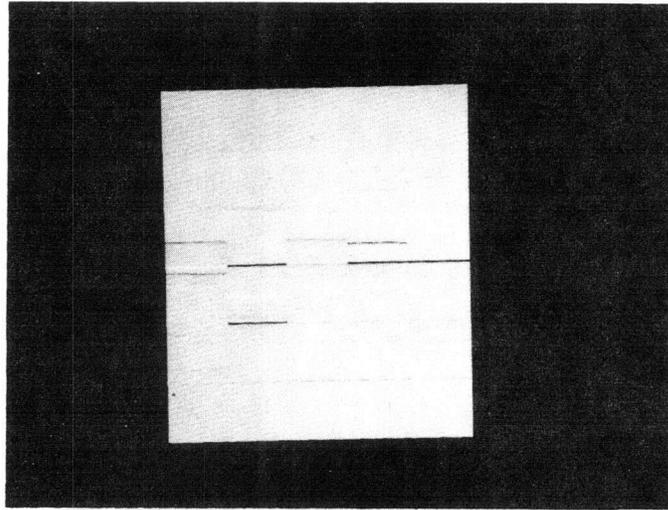


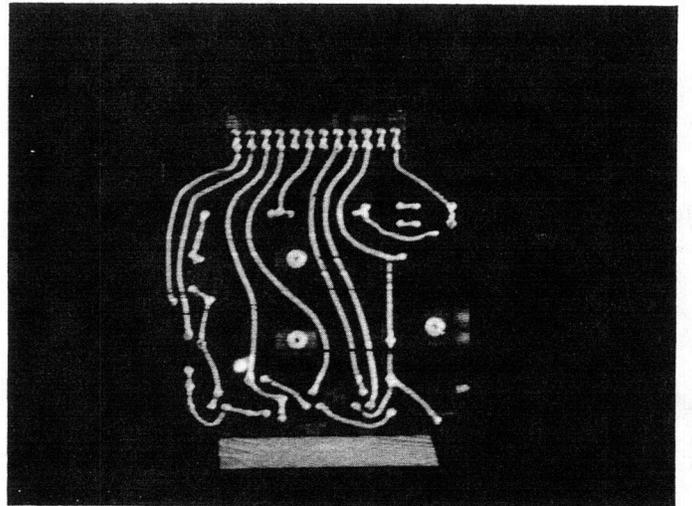
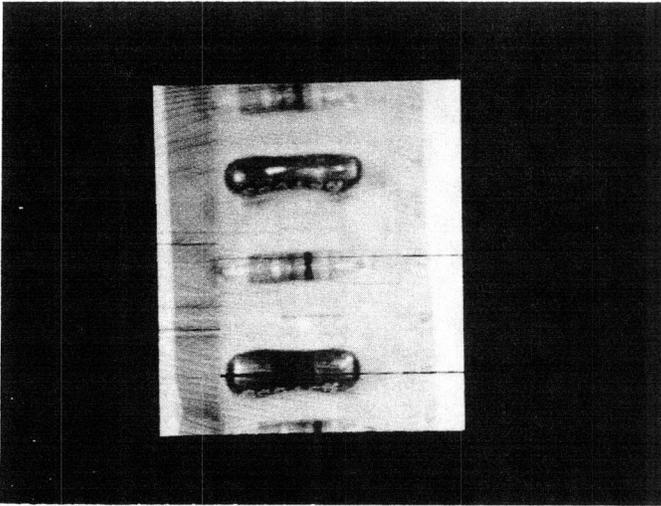
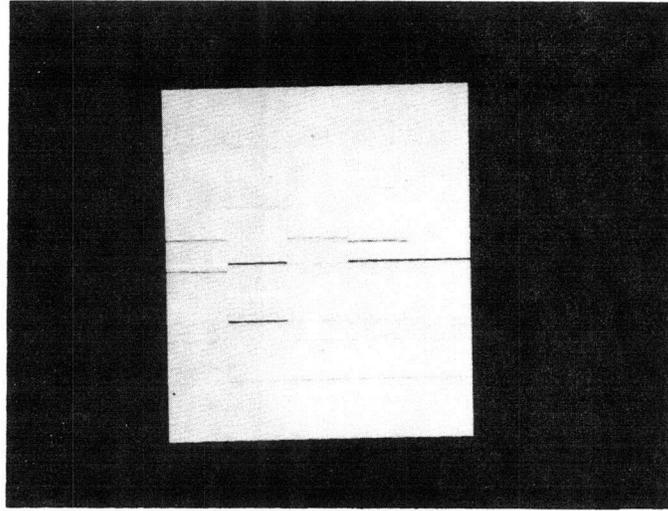


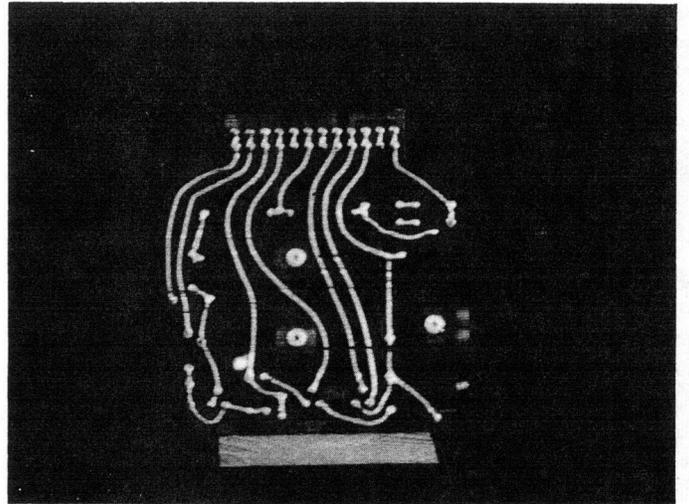
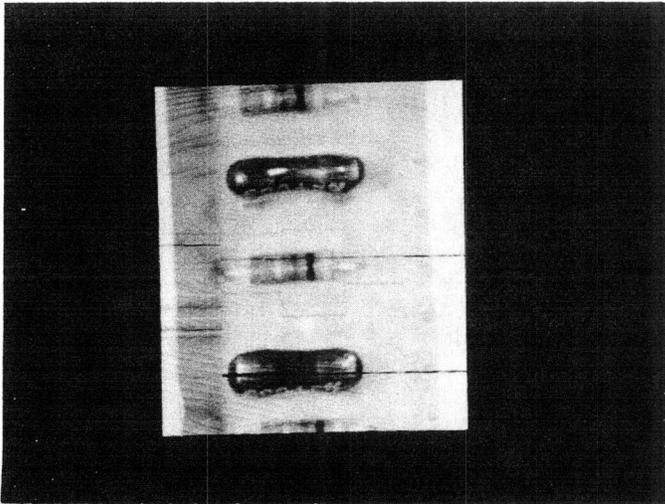
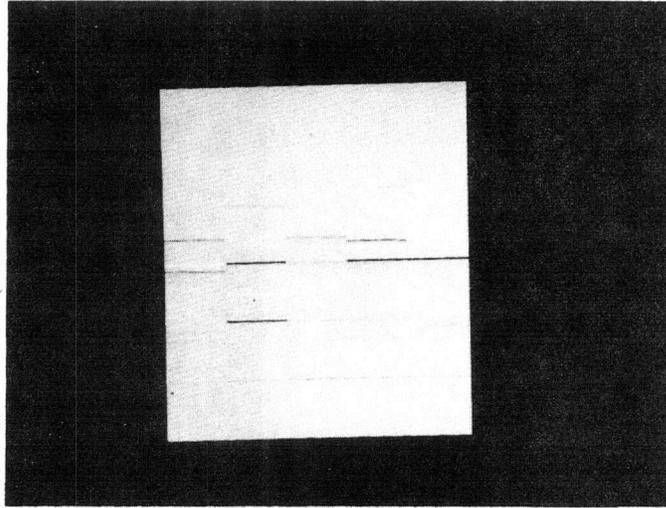
G

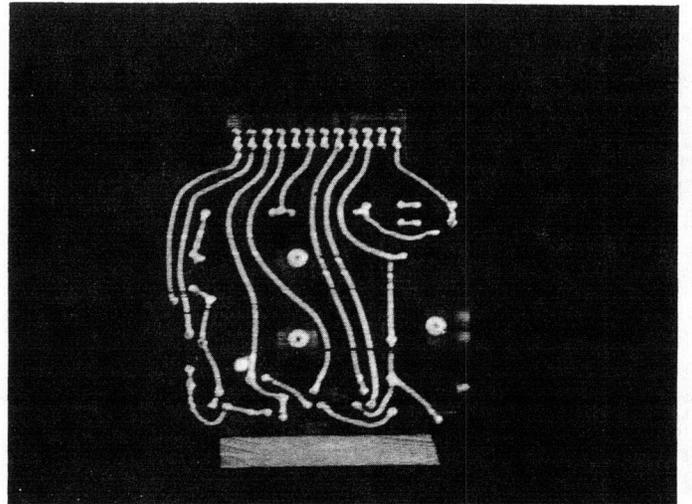
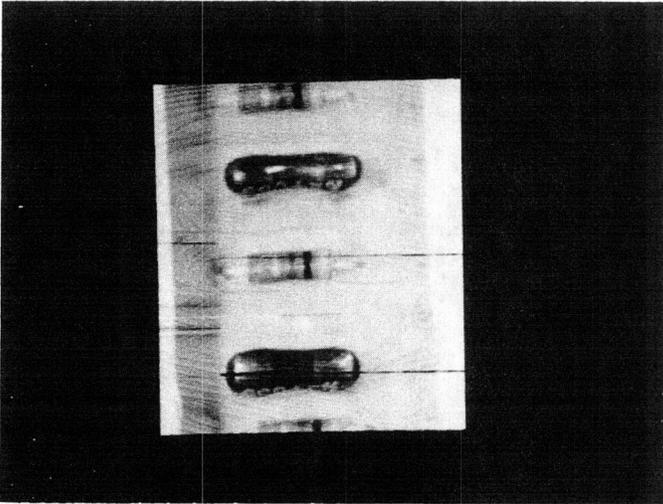
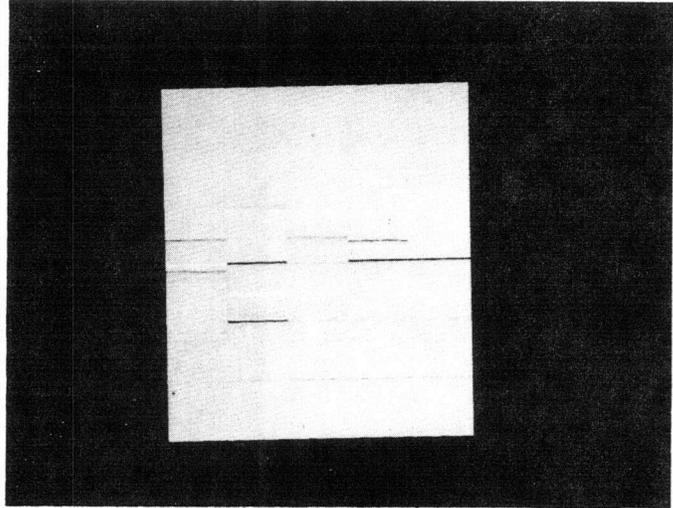


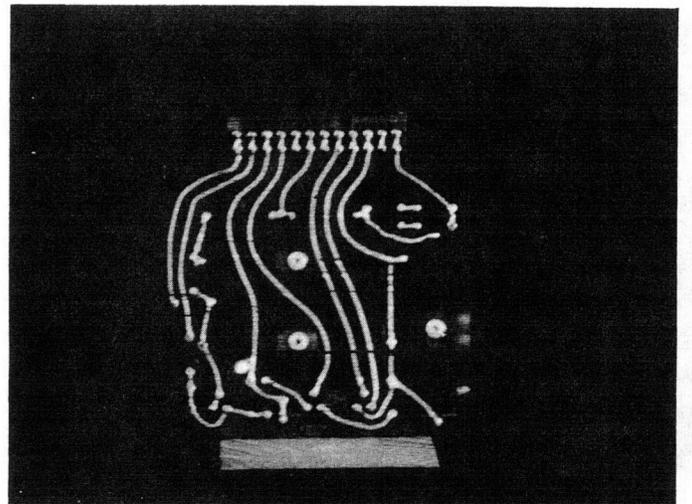
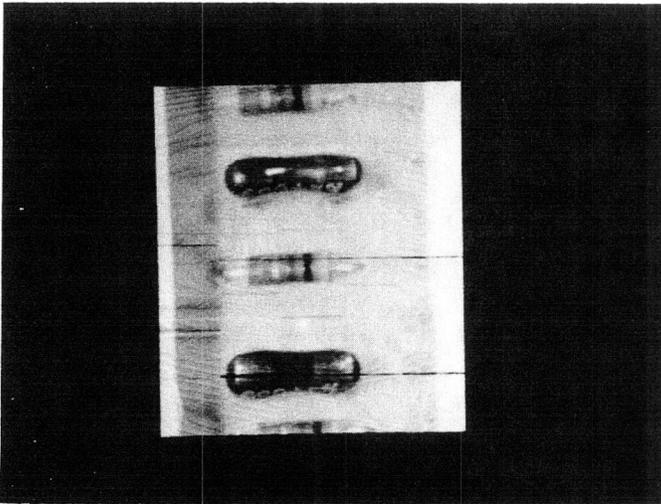
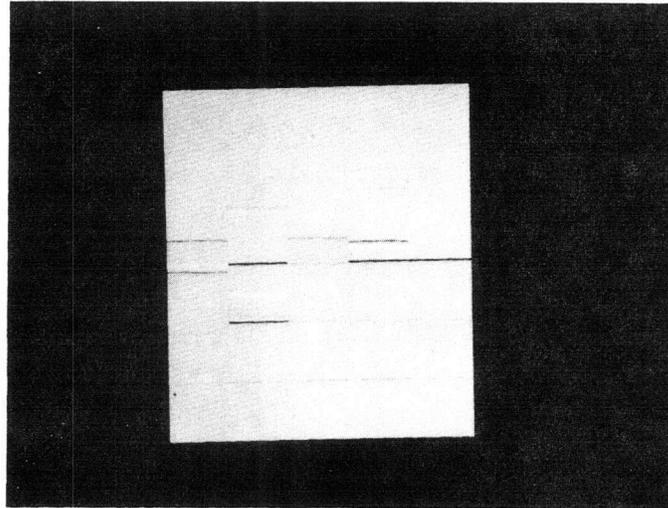


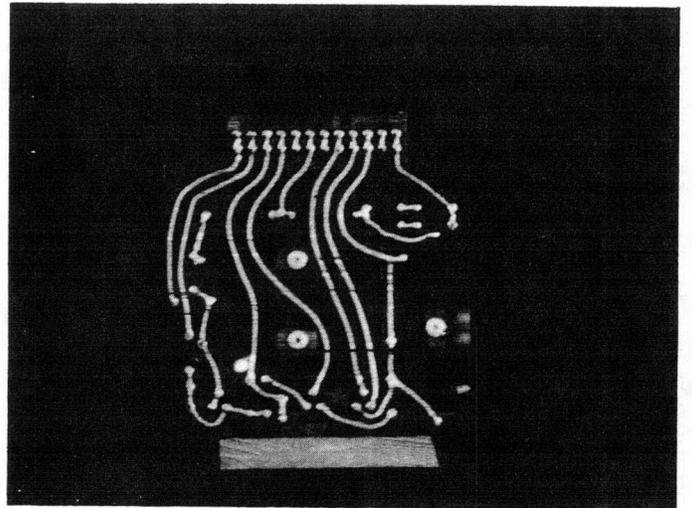
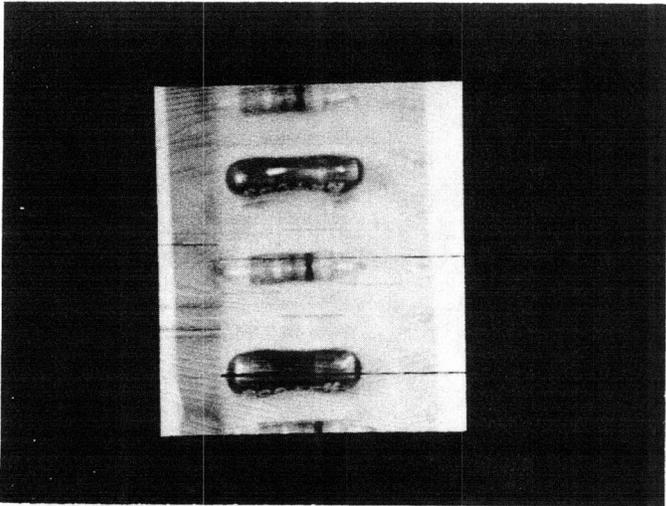
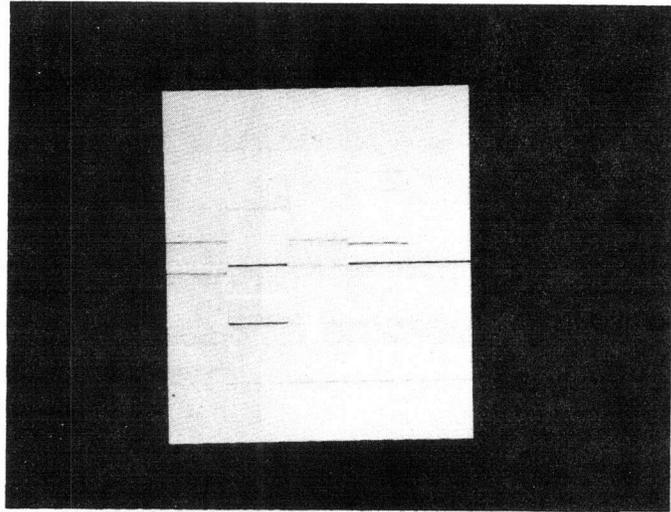


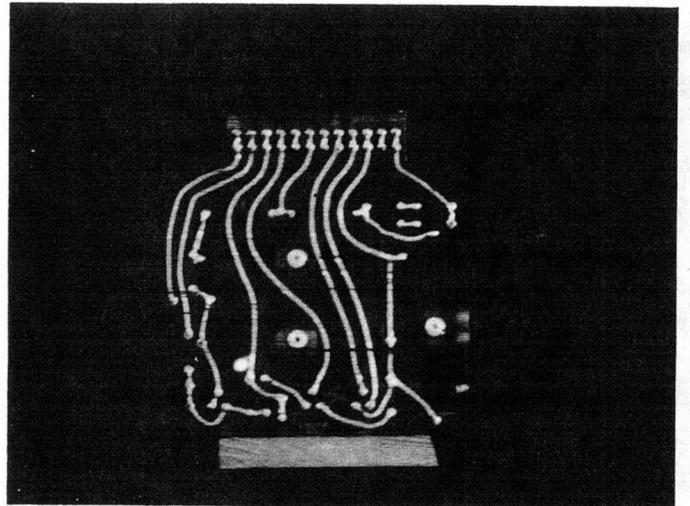
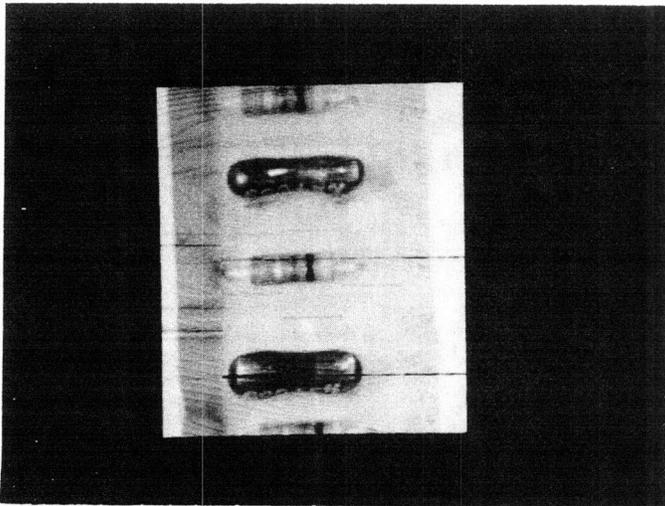
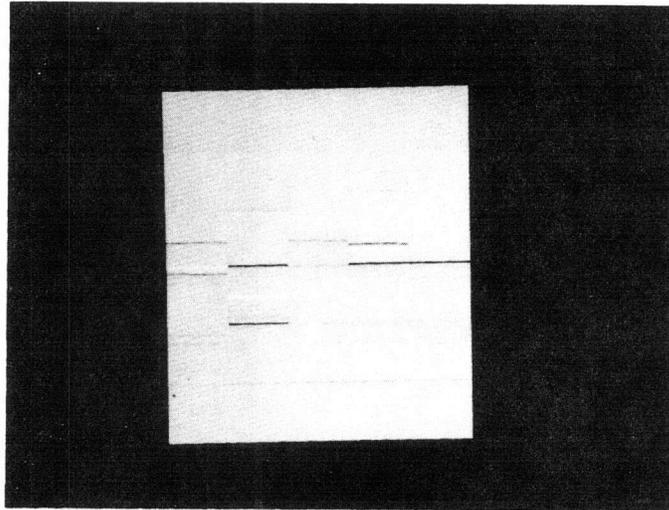


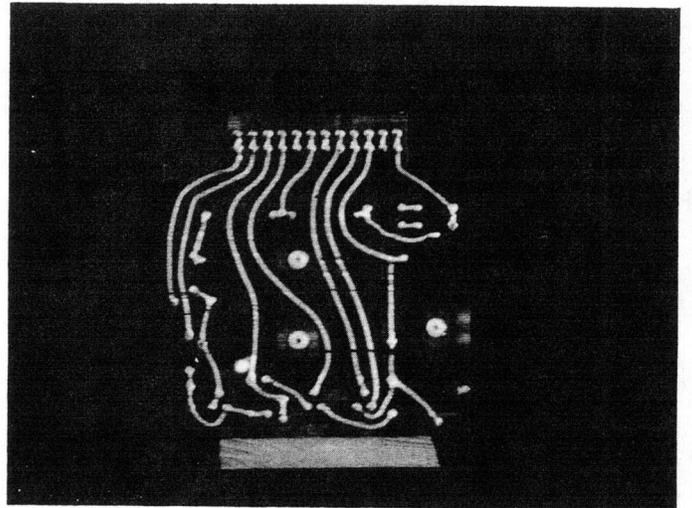
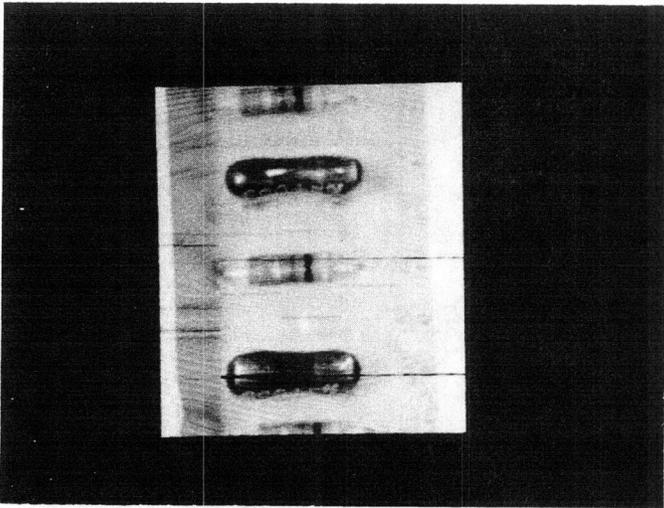
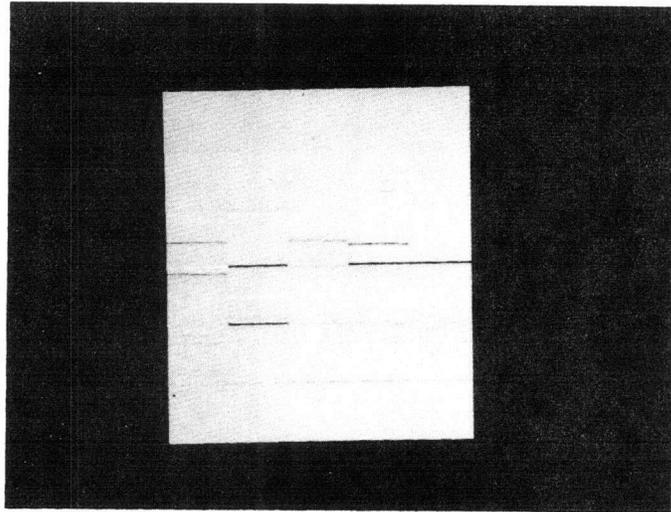


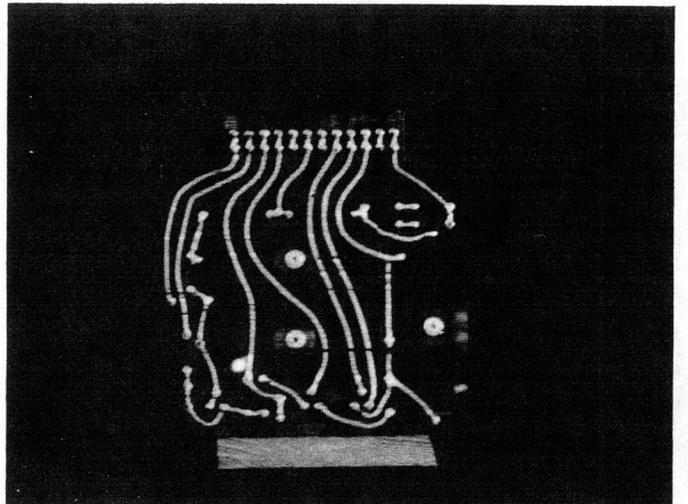
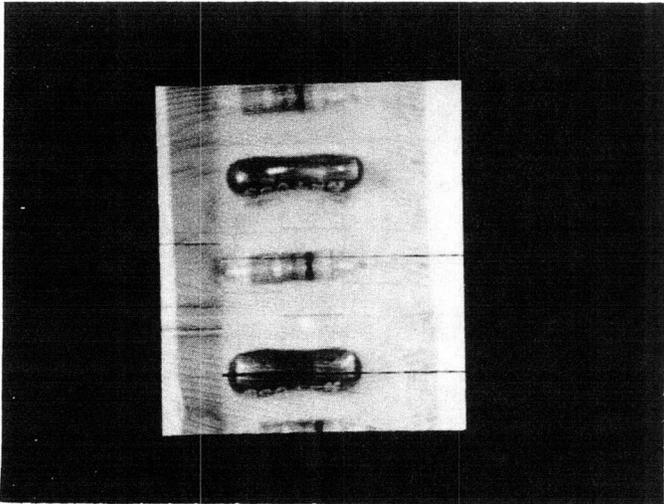
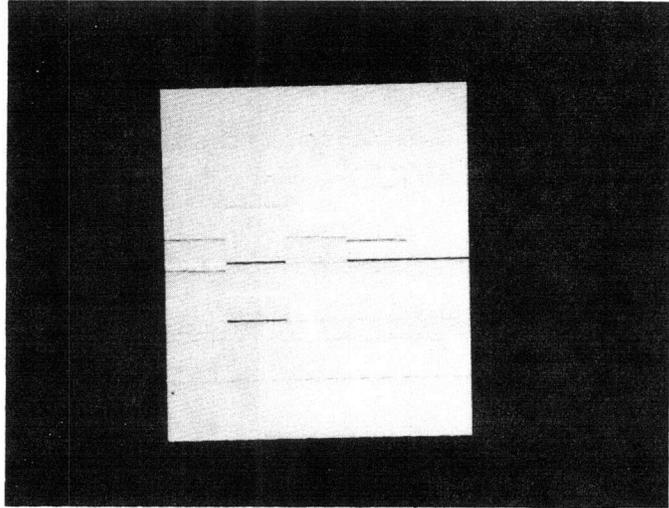












7

