

Summary

An Inexpensive Device for Pictorial Input and Output
C. R. Blair and W. W. Marshman
Department of Defense
Washington 25, D. C.

Several interesting experiments in pattern recognition by machine have been performed by manually inserting digital representations of the patterns into a computer. However, this process is tedious and admits the possibility of the experimenter inadvertently assisting the machine by supplying it with already simplified patterns. These difficulties can be circumvented by automatic pattern digitalization; we decided, therefore, to add this facility to our composer.

The basic problem of constructing a device for pictorial input and output is raster generation. For once one has devised a means of illuminating any selected point of an array, he can use it as an input device by placing the pattern on a transparency between the raster and a photocell connected to the computer input. Or he can use it for output by replacing the transparency and photocell with a camera.

Two types of raster generation are in use. The first, an adaptation of the beam-positioning techniques of electrostatic storage systems, offers the valuable advantage of high-speed random access to any point on the raster. Unfortunately, this technique requires registers at a cost too great for the funds which were available for this research. A facsimile-type scanner, on the other hand, provides a low-cost form of raster generation; it, however, is limited to sequential access to the points on the raster. Since available facsimile scanners are also quite slow, it is necessary to store the complete pattern, a circumstance far more severe than it superficially appears to be. For example, we have found that in order to extract the details - which on occasion (consider the tail of a **Q**) can be crucial - one must divide a 2 inch pattern into a 256 x 256 point array. Even if the gray scale representation of each of these points is

limited to a single bit, this pattern would occupy two-thirds of our computer's storage (4,096 words of 24-bits each). Adding even one more bit to the gray scale exceeds the entire memory capacity. And we had planned to experiment with 64-level gray scales. We decided, therefore, to design a compromise device.

An oscilloscope, the ever-present fixture of a computer installation, suggested an answer. Clearly, it could illuminate any point on the face of its cathode ray tube. The only problem remaining was to devise some means of placing this ability under program control. The oscilloscope's use of a triggered sawtooth to generate its x axis offered the clue. By arranging to have the z axis (intensity) normally turned off, any point along the x axis could be illuminated by momentarily increasing the intensity at a determinate time after the x axis is triggered. An identical sawtooth generator for the y axis permits extending the technique to two-dimensions. This facility is available as a plug-in unit in the Tektronix Type-535 oscilloscope which we used. Once facilities have been provided for triggering the x , y , and z axes of the oscilloscope under program control, the remainder of the process is straightforward. Some unused external function codes on our computer (a Sperry-Rand X-308) were used to initiate these triggers, although any 3 independent sources of pulses with a relatively fast rise time and accurate timing can be used.

Before plotting a point, both the x and the y sawtooth generators must be at their rest condition and the (turned-off) electron beam is therefore in the lower left corner of the array. When a single generator is triggered, the beam travels along the axis; when the second generator is triggered, the beam proceeds upward at a 45° angle. Thus, turning on the two sawtooth generators at the proper relative time can cause the electron beam to pass over any selected point on the face of the cathode ray tube. In order to illuminate any point on the line traced by the beam, it is only necessary momentarily to turn on the z axis (illuminate the spot) at the time that the beam passes over it. In our early experiments, we triggered both sweeps simultaneously to plot

points along the main diagonal or turned on the x or y sweeps first, depending on whether the point desired was below or above this diagonal. Unfortunately, timing difficulties caused a slight discontinuity along the diagonal of the raster. Although this dislocation could be made extremely small by shifting the plotted axes relative to the true axes, it required an excessive amount of program space. We found it simpler to trigger the x axis first and to reposition the resting point of the electron beam so that only the lower righthand quadrant of the original array was used.

For input, a transparency holder is positioned over the face of the cathode-ray tube and covered by a light-tight box which contains a photocell. Under program control the raster can illuminate any point of this transparency and the photocell output voltage indicates how well the transparency transmits light at that point. Since we already had an analog voltage-to-digital converter on the X-308, we chose to feed the output of the photocell directly into this converter. In this fashion we are able to measure the transmittance of the transparency to a precision of 64 levels. However, a much simpler device could be used for conveying a one-bit digitalization of the photocell output to the input register on a computer which does not have such a converter.

In order to produce a picture, we attach an ordinary oscilloscope camera and, with the shutter open, plot the desired points.

Of course, it is not necessary for a programmer using this facility to adopt our unusual coordinate system or to maintain the accurate timing required. He uses a closed sub-routine which, when supplied with x and y coordinates, plots the point on the oscilloscope. When the photocell is attached, the routine supply a 6-bit number representing the transmittance of the transparency at that point.

Although it is possible to plot a family of points on each transit of the electron beam, the normal usage is to produce only one point per transit. In our implementation this action requires ten milliseconds; although the time could be reduced at the expense of the excellent rectilinearity we obtain at this speed.

The extreme simplicity and low cost of this device is made possible by using the computer to perform its "logic", a technique that is too often overlooked in modern systems design. Of course, this arrangement too has its costs; in this and most similar applications, the low initial cost is offset by the expense of the additional machine time used and, in many cases, this "hidden" cost would be prohibitive. However, we believe that a large number of infrequently used computer ancillaries could profitably be constructed by using this principle.