

A Soft-edged Character Set and its Derivation

A. J. Wilkes and N. E. Wiseman

University of Cambridge, Computer Laboratory, Corn Exchange Street, Cambridge CB2 3QG, UK

Characters are normally displayed on raster display devices in the form of dot arrays. Humble VDUs often use 6×10 arrays spaced on an 8×11 grid of positions. Each character is then defined by 60 bits and the entire character set is stored in a read only memory of, say, 96 (characters) \times 6 (bits wide) \times 10 (bits high). This is 5760 bits all together and accounts for only a tiny proportion of the component cost of the unit electronics. In some of the 'better' terminals more characters may be provided. Some character sets may be held in random access memory and thus be loaded dynamically from a user's definitions, but it remains the case that the character generator logic is relatively small compared with everything else. Indeed, this is increasingly true with more and more upmarket displays—most of the increased function being devoted to picture storage and manipulation. A binary dot array is not in fact the best way to render a letterform on a CRT display. The dottiness produces a pretty crude and unsatisfactory image which is ugly and often hard to read. Raster displays do, it is true, tend always to be seen showing dotty and ugly pictures but diagrams are not intended for 'reading' in quite the same manner as running text. Fortunately, this dottiness is not intrinsic and solutions are known for arbitrary images, including text.^{1,2} Quite stunning improvements can be achieved by simple methods which can be cheaply implemented for character displays. This paper discusses the techniques and presents a complete alphabet definition as an example of the method in use.

INTRODUCTION

A picture generated for output on a device of limited resolution ought to be processed to exploit the device characteristics as well as possible so that the distortion and degradation are controlled and minimized. Except for stylized pictures tailored to the display properties, the only freedom we have is in the way the display is controlled to reproduce the given picture. The computation for a raster display is in fact a sampling of the image at a spatial frequency given by the pixel spacing. Consider the edge of some feature in the picture where differently shaded polygons meet (see Fig. 1).




Figure 1. A feature boundary intersecting the pixel grid.

This edge intersects the pixel grid, which is the spatial sampling frequency, and at each sample (pixel) we must compute an appropriate shade. Except when the edge aligns perfectly with the grid, this shade will be neither that of one polygon or the other but an intermediate value fixed by the amount of each which falls within the pixel boundary. Suppose that the edge is black on one side and white on the other, and that 4 bits are available to specify a shade encoded as a binary value between zero and 15. The situation shown above might then give rise to the following pixel values (see Fig. 2).

This use of grey scale to form what is intended as




Figure 2. (a) Grey values in vicinity of boundary. (b) Resulting image.

simply a black-white edge works extremely well when the pixel size is reasonably small. The subjective effect is in fact of a hard straight edge rather than of the soft defocussed edge which one might expect. The reason for this has to do with the way in which the brain appears to seek and accentuate line features using global clues as well as local detail. If no grey scale is available the actual pixel values will of course be zero when the wanted value is below eight, and one otherwise which produces the jagged edge which is the familiar and objectionable effect we are trying to rid ourselves of here (see Fig. 3).

The jaggedness obviously will vary in detail as the edge moves over the pixel grid. A given edge may in fact show itself in many different forms, called aliases, and a slowly moving image produces edges which shimmer as a result of the succession of different aliases which the




Figure 3. (a) Binary values in vicinity of boundary. (b) Resulting image.

movement induces. The use of grey scale in the manner suggested above is called anti-aliasing; the grey values may be computed in several different ways. The basic idea comes from sampling theory, which states that with a sampling frequency of f , accurate reconstruction may be achieved for images containing no components above $f/2$. Put simply, this means that if an image is filtered so that such high frequencies are removed, then the jaggies will go away. All methods for anti-aliasing amount to filtering the image in this way, but the methods differ in their accuracy and the amount of computing required. Quite simple, quick algorithms are often good enough, although rendering movement satisfactorily is difficult. For the present purpose a simple method will be used, based on pixel averaging. The letterforms will be specified to a high resolution so that they can be mapped onto the pixel grid by amalgamating squares. The total of black squares per pixel is then counted and used to emit a corresponding grey value for each pixel.




Figure 4. Super-sampling a pixel to determine its grey value.

Shown in Fig. 4 is a single pixel divided into a binary array of dots, presumed fine enough to allow a count of dots to be a sufficiently accurate indicator of the grey value needed. Letterforms when designed will usually be to an extremely high resolution so this condition is easily met.

In summary, the derivation of an alphabet to a low spatial resolution from binary high resolution masters can use grey values to restore some of the lost detail and fool the eye into hardening and smoothing the image contours. Shown in Figs 5-7 are some examples of the sort of results we have obtained starting with a particular alphabet and deriving others from it.




Figure 5. Enlarged letters, 20 pixels high, 2 bits per pixel.



Figure 6. Enlarged letters, 56 pixels high, 4 bits per pixel.




Figure 7. Normal sized letters, 56 pixels high, 4 bits per pixel.

IMPLEMENTATION

Experiments in letter design and text processing for print quality applications have been going on in this department for some years.³⁻⁵ Alphabet definitions associated with this work were available to us in the form of high resolution polygonal outlines. Each letter is formed from the superimposition of a few such polygons, comprising up to a hundred edges or so in a coordinate space delimited by 16 bit integers. A program to scale and scan convert these letterforms into a run-encoded form for writing text to precision raster plotters and film setters was used to preprocess the letters for the use described here. Thus our input was a run-encoded binary array for each letter on a grid roughly 700×700 . The pixel averaging algorithm we used is shown below.

```
$( LET pixels = VEC width - 1
                                // Accumulated values,
                                // which start off at zero
FOR i = 0 TO width - 1 DO pixels[i] := 0
FOR line = 1 TO N              // Read in N lines of pixel
DO $(                           // runs and accumulate them
    UNTIL end_of_line
    DO $( LET first = READN ()
          // Read in a span of blacks
          LET last = READN ()
          FOR x = first TO last
          DO pixels[(x/N) + ] := 1
    )
  )
// Write out a line
FOR x = 0 TO width - 1 DO write( pixels[x] )
newline ()
$) REPEATUNTIL end_of_data
```

The output consists of rows of *width* pixel values, the whole array being terminated by a value of -1. In practice; some compression is used: leading and trailing zeroes are suppressed and each row is preceded by an offset and the number of non zero data values on it.

The pixel values will range between 0 and N^2 and will need to be mapped to the actual set of grey values which the output device can handle. The choice of mapping is up to us. We can make use of the opportunity to tune up the performance of the overall method to compensate for brightness non-linearities in the display, filtering faults and the human vision process. If the output device has more grey values than the filter emits, then the mapping can be done by a table of outvalues indexed by the invalues. If, on the other hand, the output device has fewer grey values than the filter emits—as was the case for our experiments—the mapping can be done by a binary search in a table of invalues. The search code is:

```

outvalue := 0
increment :=  $N^2/2$ 
UNTIL (increment = 0)
DO $(
    IF (invalue > mactable!(outvalue +
        increment))
    THEN outvalue +:= increment
    increment := increment/2
$)

```

The vector *mactable* records the *invalues* at breakpoints in the mapping. The *outvalue* deduced by the algorithm is the index to the next lower *invalue* to the one given. It evidently requires that the mapping is monotonic increasing. The filtered and remapped data for each letterform can now be stored in a pixel array, say *width* pixels wide, *height* pixels high and *depth* bits deep. An alphabet in this form for *width* = 19, *height* = 21 and *depth* = 4 is reproduced in the Appendix. The mapping chosen for this example is linear—that is the *invalues* were arranged in 16 equal intervals, each corresponding with a particular *outvalue*.

REMAINDERING

In the latter case, where there are more *invalues* than *outvalues*, a cumulative quantization error is introduced, causing the integrated grey level of the resulting picture to be too low by an amount determined by the image




Figure 8. Enlarged letters, 30 pixels high, 1 bit per pixel, no remaindering.

details. Each consultation of *mactable* gives an error (or remainder) of

$$\text{invalue} - \text{mactable!outvalue}$$

which is positive or zero. The cumulative effects can be



Figure 9. Enlarged letters, 30 pixels high, 1 bit per pixel, remaindered.

minimized by distributing the error over neighbouring pixels, rather than simply throwing it away. The simple technique of adding it to the pixel value to its right works quite well, as Figs 8 and 9 show.

Remaindering in both coordinates is theoretically better but probably unnecessary when several levels of grey are available for the output. As might be expected, the effect is most noticeable when the size of the error terms is large: that is, there are few grey levels available in the output. With 16 levels, we could not easily distinguish the effects of remaindering for the alphabets we wished to display on a Sigma 5684 terminal.

USING THE DATA

The use of anti-aliased characters on displays which have a grey scale capability is relatively straightforward. A line of text is scanned *height* times, constructing and emitting on each scan the corresponding pixels for each letter required in turn. The basis of the method is shown below. A vector *txt* holds the line of text being processed, with its first word indicating the number of characters. Each character is represented by the address where its pixel array starts; a pixel array is a vector of *height* rows each held as a vector of *width* pixels.

```

FOR i = 1 TO height
$( scanptr := 1
  FOR j = 1 TO txt!0
    $( letter := text ! j
      FOR k = 1 TO width
        DO $( scanline!scanptr := (letter!i) ! k
            scanptr +:= 1
        $)
      $)
    scanline!0 := scanptr
    emit( scanline )
  $)
$)

```

A piece of logic to do the same thing is simple and would be a quite insignificant addition to the hardware of a character based display. The pixel arrays are held in ROM, *depth* bits deep instead of one, and the *depth* bits shifted to the video generator pass through a Digital to Analogue Converter (DAC) en route.

The row counter (*height* bits) and character code address the ROM in the usual way causing the emission of *width* × *depth* bits (instead of *width*). These pass into *depth* shift registers whose outputs are connected to the

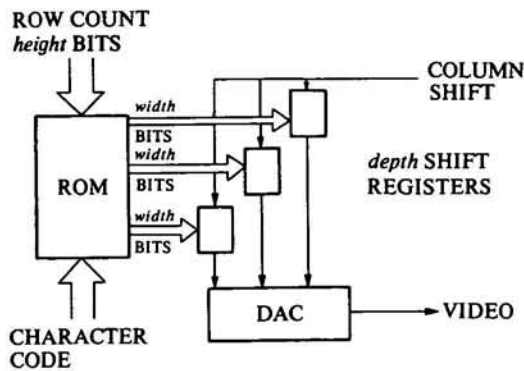


Figure 10. Basic logic for a soft-edged character generator.

DAC which emits video. No timings are altered and the circuits are not expensive.

INTERLACING

With traditional simple VDUs, the use of interlacing does not greatly increase the overall legibility of the character set, which will usually have been designed on a simple dot matrix pattern. Doubling the resolution allows a few refinements, but no major improvements are possible: the image is still locked to the pixel grid. When using a grey level display, however, advantage can be taken of any additional spatial resolution: the increased number of pixels can be immediately taken advantage of (no laborious tweaking by hand is necessary), and will be reflected directly in the quality of the final image. The example character set in the Appendix was sized to allow display of 24 lines of good quality text on a 512 line interlaced screen.

FURTHER USES

This exchange of grey values for spatial resolution in images as perceived is a reversible effect—that is, we can also use binary pictures to create a grey scale effect by losing spatial resolution (this is the means by which continuous tone images are turned into halftone images for printing by ink transfer). It seems very likely that a digital film setter such as is used to make page images for printing could exploit soft edging in the way described in this paper to remove the bumpiness in letters of which the industry constantly complains. The additional apparent spatial resolution which soft edging produces should still be there after the image is thresholded by the film emulsion, provided that image integration prior to exposure is well chosen. We hope to be able to verify this soon.

CONCLUSIONS

We have presented a simple technique for anti-aliasing characters and demonstrated the benefits in picture quality which can be easily achieved. As better displays become available, we think some such technique should be implemented in hardware. Suitable enhancements to handle variable widths and heights are not difficult to devise and, indeed, the example character set given below is intended for use with variable pitch.

By fixing the size in relation to the display resolution and accepting less freedom in the letterforms a soft edged character set could be designed which would reproduce even better than the examples shown here. Graphics arts designers might not find the restrictions of form fully acceptable, but we believe that the raster graphics medium deserves their attention.

REFERENCES

1. F. C. Crow, The aliasing problem in computer generated shaded images. *Communications of the ACM* 20, (1977), pp. 799–805.
2. J. E. Warnock, The display of characters using grey level sample arrays, *Xerox Palo Alto Research Center Research Note CSL-80-6*, (1980).
3. A. M. Pringle, P. Robinson and N. E. Wiseman, Aspects of quality in the design and production of text. *ACM SIGGRAPH* 13, (1979), pp. 63–70.
4. N. E. Wiseman, C. I. C. Campbell and J. Harradine, On making graphic arts quality output by computer. *Computer Journal* 21, (1978), pp. 2–6.
5. N. E. Wiseman, Computer designed letters. *Information Design Journal* 1, (1980), pp. 218–222.

Received May 1981

© Heyden & Son Ltd, 1982

APPENDIX

The example is based on an alphabet designed by David Kindersley in digital form using the ELF design system. The data is presented as rows of hexadecimal numbers, each one describing one line of a character. The data for each character is prefixed by \$<character> <ASCII code in octal>. The sizes have been normalized to a total character height of 21 units, and a nominal width of 19. Figure 11 shows a screen reproduction of the alphabet drawn at actual size.

Figure 11. A screen reproduction of the alphabet drawn at actual size.



[illegible]

Downloaded from <https://academic.oup.com/compl/article/25/1/140/527310> by guest on 09 April 2024

THE COMPUTER JOURNAL, VOL. 25, NO. 1, 1982 145

[illegible]

Downloaded from <https://academic.oup.com/inl/article/25/1/140/527310> by guest on 09 April 2024

THE COMPUTER JOURNAL, VOL. 25, NO. 1, 1982 147