

Fonts for Digital Halftones

Donald E. Knuth
Stanford University

Small pictures can be “typeset” on raster devices in a way that simulates the screens used to print fine books on photography. The purpose of this note is to discuss some experiments in which METAFONT has created fonts from which halftones can be generated easily on laser printers. High levels of quality are not possible at low resolution, and large pictures will overflow T_EX’s memory at high resolution; yet these fonts have proved to be useful in several applications, and their design involves a number of interesting issues.

I began this investigation several years ago when about a dozen of Stanford’s grad students were working on a project to create “high-tech self-portraits” [see Ramsey Haddad and Donald E. Knuth, “A programming and problem-solving seminar,” *Stanford Computer Science Report 1055* (Stanford, California, June 1985), pp. 88–103]. The students were manipulating digitized graphic images in many ingenious ways, but Stanford had no output devices by which the computed images could be converted to hardcopy. Therefore I decided to create a font by which halftones could be produced using T_EX.

Such a font is necessarily device-dependent. For example, a laser printer with 300 pixels per inch cannot mimic the behavior of another with 240 pixels per inch, if we are trying to control the patterns of pixels. I decided to use our 300-per-inch Imagen laserprinter because it gave better control over pixel quality than any other machine we had.

It seemed best at first to design a font whose “characters” were tiny 8×8 squares of pixels. The idea was to have 65 characters for 65 different levels of brightness: For $0 \leq k \leq 64$ there would be one character with exactly k black pixels and $64 - k$ white pixels.

Indeed, it seemed best to find some permutation p of the 64 pixels in an 8×8 square so that the black pixels of character k would be p_0, p_1, \dots, p_{k-1} . My first instinct was to try to keep positions p_0, p_1, p_2, \dots as far apart from each other as possible. So my first METAFONT program painted pixels black by ordering the positions as follows:

45	29	34	18	46	30	33	17
13	61	2	50	14	62	1	49
39	23	40	24	36	20	43	27
7	55	8	56	4	52	11	59
47	31	32	16	44	28	35	19
15	63	0	48	12	60	3	51
37	21	42	26	38	22	41	25
5	53	10	58	6	54	9	57

[This is essentially the “ordered dither” matrix of B. E. Bayer; see the survey paper by Jarvis, Judice, and Ninke in *Computer Graphics and Image Processing* 5 (1976), 22–27.]

It turns out to be easy to create such a font with METAFONT:

```
% halftone font with 65 levels of gray, characters "0" (white) to "p" (black)
pair p[]; % the pixels in order (first p0 becomes black, then p1, etc.)
pair d[]; d[0]=(0,0); d[1]=(1,1); d[2]=(0,1); d[3]=(1,0); % dither control
def wrap(expr z)=(xpart z mod 8,ypart z mod 8) enddef;
for i=0 upto 3: for j=0 upto 3: for k=0 upto 3:
  p[16i+4j+k]=wrap(4d[k]+2d[j]+d[i]+(2,2)); endfor endfor endfor
w#:=8/pt; % that's 8 pixels
font_quad:=w#; designsize:=8w#;
picture prevchar; prevchar=nullpicture; % the pixels blackened so far
for i=0 upto 64:
  beginchar(i+ASCII"0",w#,w#,0); currentpicture:=prevchar;
  if i>0: addto currentpicture also unitpixel shifted p[i-1]; fi
  prevchar:=currentpicture; endchar;
endfor
```

This file was called `dt.mf`; I used it to make a font called 'dt300' by applying METAFONT in the usual way to the following file `dt300.mf`:

```
% Halftone font for Imagen, dithered
mode_setup;
if (pixels_per_inch<>300) or (mag<>1):
  errmessage "Sorry, this font is only for resolution 300";
  errmessage "Abort the run now or you'll clobber the TFM file";
  forever: endfor
else: input dt fi
end.
```


(The purpose of `dt300.mf` is to enforce the device-dependence of this font.)

It's fairly easy to typeset pictures with `dt300` if you input the following macro file `hf65.tex` in a \TeX document:

```
\font\halftone=dt300 % for halftones on the Imagen 300
\chardef\other=12
\def\beginhalftone{\vbox\bgroup\offinterlineskip\halftone
\catcode'\=\other \catcode'\^=\other \catcode'\_=\other
\catcode'\.=\active \starhalftone}
{\catcode'\.=\active \catcode'\/=0 \catcode'\\=\other
/gdef/starhalftone#1\endhalftone{/let.=/endhalftoneline
/beginhalftoneline#1/endhalftone}}
\def\beginhalftoneline{\hbox\bgroup\ignorespaces}
\def\endhalftoneline{\egroup\beginhalftoneline}
\def\endhalftone{\egroup\setbox0=\lastbox\unskip\egroup}
% Example of use:
% \beginhalftone
% chars for top line of picture.
% chars for second line of picture.
% ...
% chars for bottom line of picture.
% \endhalftone
```

(These macros are a bit tricky because `\` is one of the legal characters in `dt300`; we must make backslashes revert temporarily to the status of ordinary symbols.)


Unfortunately, the results with `dt300` weren't very good. For example, here are three typical pictures, shown full size as they came off the machine:*



The squareness of the pixels is much too prominent.

* Asterisks are used throughout this paper to denote places where output from the 300-pixels-per-inch Imagen printer has been pasted in. Elsewhere, the typesetting is by an APS Micro-5, which has a resolution of about 723 pixels per inch.


Moreover, the laser printer does strange things when it is given pixel patterns like those in dt300:*



Although character k has more black pixels than character $k - 1$, the characters do not increase their darkness monotonically! Character 6 seems darker than character 7; this is an optical illusion. Character 32 is darker than many of the characters that follow, and in this case the effect is not illusory: Examination with a magnifying glass shows that the machine deposits its toner in a very curious fashion.

Another defect of this approach is that most of the characters are quite dark; 50% density is reached already at about character number 16. Hence dt300 overemphasizes light tones.


My next attempt was to look at halftone pictures in books and newspapers, in order to discover the secret of their success. Aha! These were done by making bigger and bigger black dots; in other words, the order of pixels p_0, p_1, \dots was designed to keep black pixels *close together* instead of far apart. Also, the dots usually appear in a grid that has been rotated 45° , since human eyes don't notice the dottiness at this angle as much as they do when a grid is rectilinear. Therefore I decided to blacken pixels in the following order:



Here I decided not to stick to an 8×8 square; this nonsquare set of pixel positions still “tiles” the plane in Escher-like fashion, if we replicate it at 8-pixel intervals. The characters are considered to be 8 pixels wide and 8 pixels tall, as before, but they are no longer confined to an 8×8 bounding box. The reference point is the lower left corner of position 24.


The matrix above is actually better than the one I first came up with, but I've forgotten what that one was. John Hobby took a look at mine and suggested this alternative, because he wanted the pattern of *black* pixels in character k to be essentially the same as the pattern of *white* pixels in character $64 - k$. (Commercial halftone schemes start with small black dots on a white background; then the dots grow until they form a checkerboard of black and white; then the white dots begin to shrink into their black background.) The matrix above has this symmetry property, because the sum of the entries in positions (i, j) and $(i, j + 4)$ is 63 for all i and j , if you consider “wraparound” by computing indices modulo 8.

John and I used this new ordering of pixel positions to make a font called dot300, analogous to dt300. It has the following gray levels:*



Now we have a pleasantly uniform gradation, except for an inevitable anomaly between characters 62 and 63. The density reaches 50% somewhere around character number 45, and we can compensate for this by preprocessing the data to be printed.

The three images that were displayed with dt300 above look like this when dot300 is used:*



My students were able to use dot300 successfully, so I stopped working on halftones and resumed my normal activities.

However, I realized later that dot300 can easily be improved, because each of its characters is made up of two dots that are about the same size. There's no reason why the dots of a halftone image need to be paired up in such a way. With just a bit more work, we can typeset each dot independently!

Thus, I made a font hf300 with just 33 characters (not 65 as before), using the matrix

				31					
			19	23	27				
			14	7	11	15	17		
		26	10	2	3	5	21	29	
30	22	6	0	1	9	25			
		18	12	8	4	13			
			24	20	16				
						28			


to control the order in which pixels are blackened. (This matrix corresponds to just one of the two dots in the larger matrix above.) The characters are still regarded as 8 pixels wide, but they are now only 4 pixels tall. When a picture is typeset, the odd-numbered rows are to be offset horizontally by 4 pixels.

Here is the METAFONT file hf.mf that was used to generate the single-dot font:


```
% halftone font with 33 levels of gray, characters "0" (white) to "P" (black)
pair p[]; % the pixels in order (first p0 becomes black, then p1, etc.)
p0=(1,1); p4=(2,0); p8=(1,0); p12=(0,0);
p16=(3,-1); p20=(2,-1); p24=(1,-1); p28=(2,-2);
transform r; r=identity rotatedaround ((1.5,1.5),90);
for i=0 step 4 until 28:
  p[i+1]=p[i] transformed r;
  p[i+3]=p[i+1] transformed r;
  p[i+2]=p[i+3] transformed r;
endfor
w#:=8/pt; % that's 8 pixels
font_quad:=w#; designsize:=8w#;
picture prevchar; prevchar=nullpicture; % the pixels blackened so far
for i=0 upto 32:
  beginchar(i+ASCII"0",w#,.5w#,0); currentpicture:=prevchar;
  if i>0: addto currentpicture also unitpixel shifted p[i-1]; fi
  prevchar:=currentpicture; endchar;
endfor
```

(There's also a file `hf300.mf`, analogous to the file `dt300.mf` above.)

Here's how the three example images look when they're rendered by font `hf300`.*



They are somewhat blurry because they were generated second-hand from data intended for square pixels; sharper results are possible if the data is expressly prepared for a 45° grid. For example, here is a sharper Mona Lisa, and an image whose dots were computed directly by mathematical formulas:*



The $\text{T}_{\text{E}}\text{X}$ macros `hf65.tex` shown above must be replaced by another set `hf33.tex` when independent dots are used:

```

\font\halftone=hf300 % for halftones on the Imagen 300, each dot independent
\chardef\other=12

\newif\ifshifted
\def\shift{\moveright.5em}
\def\beginhalftone{\vbox\bgroup\offinterlineskip\halftone
  \catcode'\.=\active\shiftedtrue\shift\hbox\bgroup}
\catcode'\.=\active \gdef.\egroup
  \ifshifted\shiftedfalse\else\shiftedtrue\shift\fi\hbox\bgroup\ignorespaces}}
\def\endhalftone{\egroup\setbox0=\lastbox\egroup}

% Example of use:
% \beginhalftone
% chars for top halfline of picture. (shifted right 4 pixels)
% chars for second halfline of picture. (not shifted right)
% chars for third halfline of picture. (shifted right 4 pixels)
% ...
% chars for bottom halfline of picture. (possibly shifted right)
% \endhalftone

```

These macros are much simpler than those of `hf65`, because the 33 ASCII characters "0" to "P" have no special meaning to plain $\text{T}_{\text{E}}\text{X}$.

We can also create an analogous font hf723 for the high-resolution APS, in which case the pictures come out looking like this:



The same T_EX macros were used, but font \halfstone was defined to be hf723 instead of hf300. Now the pictures are smaller, because the font characters are still 8 pixels wide, and the pixels have gotten smaller. At this resolution the halftones look “real,” except that they are too dark. This problem can be fixed by adjusting the densities in a preprocessing program. Also, small deficiencies in the APS’s analog-to-digital conversion hardware become apparent when such tiny characters are typeset.

What resolution is needed? It is traditional to measure the quality of a halftone screen by counting the number of dots per inch in the corresponding unrotated grid, and it’s easy to do this with a magnifying glass. The photographs in a newspaper like the *International Herald Tribune* use a 72-line screen, rotated 45°; this is approximately the resolution $50\sqrt{2}$ that we would obtain with the hf400 font on a laser printer with 400 dots per inch. (The 300-per-inch font hf300 gives a rotated screen with only $37.5\sqrt{2} \approx 53$ dots per inch.) The photographs on the book jacket of *Computers & Typesetting* have a 133-line screen, again rotated 45°; this is almost identical to the resolution of hf723. But this is not the upper limit: A book that reproduces photographs with exceptionally high quality, such as *Portraits of Success* by Carolyn Caddes (Portola Valley: Tioga Press, 1986), has a screen of about 270 lines per inch, in this case rotated 30°.

Let’s turn now to another problem: Suppose we have an image for which we want to obtain the best possible representation on a laser printer of medium resolution, because we will be using that image many times—for example, in a letterhead. In such cases it is clearly desirable to create a special font for that image alone; instead of using a general-purpose font for halftones, we’ll want to control every pixel. The desired image can then be typeset from a special-purpose font of “characters” that represent rectangular subsections of the whole.

The examples above were produced on an Imagen printer as 64 lines of 55 columns per line, with 8 pixels in each line and each column. To get an equivalent picture with every pixel selected individually, we can make a font that has, say, 80 characters, each 64 pixels tall and 44 pixels wide. By typesetting eight rows of ten characters each, we’ll have the desired image. For example, the following picture was done in that way:*





T_EX will typeset such an image if we say \monalisa after making the following definitions:

```

\font\mona=mona300[hf,dek]
\newcount\m \newcount\n
\def\monalisa{\vbox{\mona \offinterlineskip \n=0
\loop \hbox{\m=0 \loop \char\n \global\advance\n by 1
\advance\m by 1 \ifnum\m<10 \repeat}
\ifnum\n<80 \repeat}}

```

And once we have the individual pieces, we can combine them to get unusual effects:*



The font mona300 shown above was generated from a file mona.mf that began like this:

```
row(1); cols(1,5,9,13,15,17,21,24,30,32,39,46,56,62,70,
             78,86,95,103,110,118,120,127,135,142,151,159,167,175,183,
             191,198,207,215,223,230,238,246,254,263,271,279,287,295,302,
             311,318,328,334,342,350,358,366,367,375,382,383,390,392,398,
             400,405,408,414,416,421,424,430,432,439);
row(2); cols(4,7,12,20,23,28,30,37,38,40,45,48,53,61,64,
```

... and so on, until 512 rows had been specified. The parameter file mona300.mf was

```
% Mona Lisa for Imagen 300
mode_setup;
if (pixels_per_inch<>300) or (mag<>1):    ...  <error messages as before>
else: input picfont
    width:=44; height:=64; m:=8; n:=10; filename:="mona";
    do_it; fi
end.
```


and the driver file picfont.mf was

```
def do_it=
  for j=0 upto n-1: jj:=width*j; jjj:=jj+width; jjjj:=j;
  scantokens("input "&filename); endfor enddef;
string filename;
def row(expr x) =
  cc:=(x-1)div height; rr:=height-1-((x-1)mod height);
  if rr=height-1: beginchar(cc*n+jjjj,width/pt,height/pt,0); fi enddef;
def cols(text t) =
  for tt:=t: exitif tt>=jjj; if tt>=jj:
  addto currentpicture also unitpixel shifted (tt,rr); fi endfor
  if rr=0: xoffset:=-jj; endchar; fi enddef;
```

This is not very efficient, but it's interesting and it seems to work.

Ken Knowlton and Leon Harmon have shown that surprising effects are possible once a picture has been digitized [see *Computer Graphics and Image Processing 1* (1972), 1–20]. Continuing this tradition, I found that it's fun to combine the T_EX macros above with new fonts that frankly acknowledge their digital nature. One needn't always try to compete with commercial halftone screens!

For example, we can use `hf65.tex` with a 'negdot' font that makes negative images out of square dots:




The METAFONT file `negdot.mf` that generated this font is quite simple:

```
% negative pseudo-halftone font with 65 sizes of square dots
mode_setup;
w#:=2.5pt#; font_quad:=w#; designsize:=8w#;
for i=0 upto 64:
  beginchar(i+ASCII"0",w#,w#,0);
  r#:=sqrt(.9w#*(1-i/80)); define_pixels(r);
  fill unitsquare scaled r shifted(.5w,.5h);
  endchar;
endfor
end.
```

Unlike the previous fonts we have considered, this one is device-independent.

It's even possible to perceive images when each character of the halftone font has exactly the same number of black pixels. Here, for example, is what happens when the three images above are typeset with a font in which each character consists of a vertical line and a horizontal line; the lines move up and to the right as the pixel gets darker, but they retain a uniform thickness. We perceive lighter and darker features only because adjacent lines get closer together or further apart.




The METAFONT file `lines.mf` for this device-independent font is:

```
% pseudo-half-tone font with 65 lines that move right and up
mode_setup; q:=savepen;
w#:=2.5pt#; font_quad:=w#; designsize:=8w#;
for i=0 upto 64: beginchar(i+ASCII"0",w#,w#,0); pickup q;
  draw (0,h*i/64)--(w,h*i/64); draw(w*i/64,0)--(w*i/64,h); endchar;
endfor end.
```

Yet another possibility is the font produced by `angles.mf`; here each character is a single line of radius 2.5pt that rotates from horizontal to vertical as the density increases:

```
% pseudo-half-tone font with 65 radii that move counterclockwise
mode_setup; q:=savepen;
w#:=2.5pt#; font_quad:=w#; designsize:=8w#;
for i=0 upto 64: beginchar(i+ASCII"0",w#,w#,0); pickup q;
  draw ((0,0)--(w,0)) rotated (90*i/64); endchar;
endfor end.
```

The images are still amazingly easy to identify:



(We can think of a large array of dials whose hands record the local light levels.) It is amusing to view these images by tilting the page up until your eyes are almost parallel to the paper.

As a final example, let's consider a 33-character font that's designed to be used with `hf33.tex` instead of `hf65.tex`. Readers who like puzzles are invited to try to guess what this METAFONT code will do, before looking at the image of Mona Lisa that was typeset with the corresponding font. [*Hint*: The name of the METAFONT file is `hex.mf`.]

```
% pseudo-half-tone font with 33 more-or-less hexagonal patterns
mode_setup; q:=savepen;
w#:=7.5pt#; font_quad:=w#; designsize:=8w#;
for i=0 upto 32: beginchar(i+ASCII"0",w#,.5w#,0); pickup q;
  alpha:=.5-i/72; z0=(.5w,.5h);
  z1=alpha[(5/6w,.5h),z0]; z2=alpha[(2/3w,-.5h),z0];
  z0=.5[z2,z5]=.5[z3,z6]=.5[z1,z4]; x2=x6; y5=y6;
  draw z1--z2; draw .5[z1,z2]--z0;
  draw z3--z4; draw .5[z3,z4]--z0;
  draw z5--z6; draw .5[z5,z6]--z0; endchar;
endfor
end.
```

The answer to this puzzle can be seen in the illustration at the very end of this paper (following the appendices).

Appendix 1: Source data for the examples

The examples in this paper are mostly derived from the basic pixel values shown below. This data uses a convention taken from the book Digital Image Processing by Rafael C. Gonzalez and Paul Wintz (Addison-Wesley, 1977): The 32 characters 0123456789ABCDEF... represent densities from 1.0 down

Table with two columns: (Lisa) and (Lincoln). Each column contains a long string of alphanumeric characters representing pixel density data for a specific image.

Appendix 3: Transforming the pixel data

The following WEB program illustrates how to convert data like that of Appendix 1 into the form required by the fonts and macros described earlier.

1. Introduction. This program prepares 33-level halftone images for use in \TeX files. The input is assumed to be a sequence of pictures expressed in the form

```

  m  n
  <first line of pixel data, n characters long>
  ...
  <mth line of pixel data, n characters long>

```

terminated by a line that says simply '0'. The pixel data consists of the characters "0" to "9" and "A" to "V", representing 32 levels of darkness from black to white. [See Appendix 1.]

The output is the same set of pictures, expressed in a simple format used for 33-level halftones, with ASCII characters "0" to "P" representing darkness levels from white to black. The levels are adjusted to compensate for the idiosyncrasies of Canon LBP-CX laser-printing engines. Two dots are typeset for each pixel of input; hence there are $2m$ "halflines" of n -character data in the output.

2. Here's an outline of the entire Pascal program:

```

program halftones(input, output);
  label <Labels in the outer block 4>
  const <Constants in the outer block 3>
  type <Types in the outer block 5>
  var <Global variables 6>

  procedure initialize; { this procedure gets things started properly }
    var <Local variables for initialization 8>
    begin <Set initial values 7>
    end;
  begin initialize; <The main program 23>;
  end.

```

3. Each picture in the input data must contain fewer than max_m rows and max_n columns.

```

<Constants in the outer block 3> ≡
  max_m = 200; { m should be less than this }
  max_n = 200; { n should be less than this }

```

This code is used in section 2.

4. The main program has one statement label, namely *cleanup_and_terminate*.

```

define cleanup_and_terminate = 9998
define finish ≡ goto cleanup_and_terminate { do this when all the pictures have been output }
<Labels in the outer block 4> ≡
  cleanup_and_terminate;

```

This code is used in section 2.

5. The character set. We need translation tables between ASCII and the actual character set, in order to make this program portable. The standard conventions of \TeX : *The Program* are copied here, essentially verbatim.

```

define text_char ≡ char { the data type of characters in text files }
define first_text_char = 0 { ordinal number of the smallest element of text_char }
define last_text_char = 127 { ordinal number of the largest element of text_char }
<Types in the outer block 5> ≡
  ASCII_code = 0 .. 127; { seven-bit numbers }

```

This code is used in section 2.

6. (Global variables 6) \equiv

xord: **array** [*text_char*] of *ASCII_code*; { specifies conversion of input characters }
xchr: **array** [*ASCII_code*] of *text_char*; { specifies conversion of output characters }

See also sections 10, 13, 14, and 22.

This code is used in section 2.

7. (Set initial values 7) \equiv

```
xchr[40] ← " "; xchr[41] ← "!"; xchr[42] ← "\"; xchr[43] ← "#"; xchr[44] ← "$"; xchr[45] ← "%";
xchr[46] ← "&"; xchr[47] ← "'";
xchr[50] ← "("; xchr[51] ← ")"; xchr[52] ← "*"; xchr[53] ← "+"; xchr[54] ← ","; xchr[55] ← "-";
xchr[56] ← "."; xchr[57] ← "/";
xchr[60] ← "0"; xchr[61] ← "1"; xchr[62] ← "2"; xchr[63] ← "3"; xchr[64] ← "4"; xchr[65] ← "5";
xchr[66] ← "6"; xchr[67] ← "7";
xchr[70] ← "8"; xchr[71] ← "9"; xchr[72] ← ":"; xchr[73] ← ";"; xchr[74] ← "<"; xchr[75] ← "=";
xchr[76] ← ">"; xchr[77] ← "?";
xchr[100] ← "@"; xchr[101] ← "A"; xchr[102] ← "B"; xchr[103] ← "C"; xchr[104] ← "D";
xchr[105] ← "E"; xchr[106] ← "F"; xchr[107] ← "G";
xchr[110] ← "H"; xchr[111] ← "I"; xchr[112] ← "J"; xchr[113] ← "K"; xchr[114] ← "L";
xchr[115] ← "M"; xchr[116] ← "N"; xchr[117] ← "O";
xchr[120] ← "P"; xchr[121] ← "Q"; xchr[122] ← "R"; xchr[123] ← "S"; xchr[124] ← "T";
xchr[125] ← "U"; xchr[126] ← "V"; xchr[127] ← "W";
xchr[130] ← "X"; xchr[131] ← "Y"; xchr[132] ← "Z"; xchr[133] ← "["; xchr[134] ← "\";
xchr[135] ← "]"; xchr[136] ← "^"; xchr[137] ← "_";
xchr[140] ← "`"; xchr[141] ← "a"; xchr[142] ← "b"; xchr[143] ← "c"; xchr[144] ← "d";
xchr[145] ← "e"; xchr[146] ← "f"; xchr[147] ← "g";
xchr[150] ← "h"; xchr[151] ← "i"; xchr[152] ← "j"; xchr[153] ← "k"; xchr[154] ← "l";
xchr[155] ← "m"; xchr[156] ← "n"; xchr[157] ← "o";
xchr[160] ← "p"; xchr[161] ← "q"; xchr[162] ← "r"; xchr[163] ← "s"; xchr[164] ← "t";
xchr[165] ← "u"; xchr[166] ← "v"; xchr[167] ← "w";
xchr[170] ← "x"; xchr[171] ← "y"; xchr[172] ← "z"; xchr[173] ← "{"; xchr[174] ← "|";
xchr[175] ← "}" ; xchr[176] ← "~";
xchr[0] ← " "; xchr[177] ← " "; { ASCII codes 0 and 177 do not appear in text }
```

See also sections 9, 11, 15, and 17.

This code is used in section 2.

8. (Local variables for initialization 8) \equiv

i: 0 .. *last_text_char*;

This code is used in section 2.

9. (Set initial values 7) $+ \equiv$

```
for i ← 1 to 37 do xchr[i] ← " ";
for i ← first_text_char to last_text_char do xord[chr(i)] ← 177;
for i ← 1 to 176 do xord[xchr[i]] ← i;
```

10. Inputting the data. We keep the pixel values in a big global array called *v*. The variables *m* and *n* keep track of the current number of rows and columns in use.

The *dd* table contains density values assumed for the input, indexed by single-character codes.

(Global variables 6) $+ \equiv$

v: **array** [0 .. *max_m*, 0 .. *max_n*] of *real*; { pixel darknesses, from 0.0 to 1.0 }
m: *integer*; { rows 0 .. *m* + 1 of *v* should contain relevant data }
n: *integer*; { columns 0 .. *n* + 1 of *v* should contain relevant data }
dd: **array** [*text_char*] of *real*;

11. All input codes give zero density, except "0" to "9" and "A" to "V".

```

⟨Set initial values  $\tau$ ⟩ +≡
  for  $i \leftarrow \text{first\_text\_char}$  to  $\text{last\_text\_char}$  do  $dd[\text{chr}(i)] \leftarrow 0.0$ ;
  for  $i \leftarrow "0"$  to  $"9"$  do  $dd[\text{chr}(i)] \leftarrow 1.0 - (i - "0")/31.0$ ;
  for  $i \leftarrow "A"$  to  $"V"$  do  $dd[\text{chr}(i)] \leftarrow 1.0 - (i - "A" + 10)/31.0$ ;

```

12. The process of inputting pixel values is quite simple. We terminate the program if anomalous values of m and n occur. Boundary values are added at the top, left, right, and bottom in order to provide "padding" that will be convenient in the pixel transformation process. Each boundary value is equal to one of its adjacent neighbors.

```

⟨Input a picture, or terminate the program 12⟩ ≡
  read( $m$ ); if ( $m \leq 0$ )  $\vee$  ( $m \geq \text{max\_m}$ ) then finish;
  read_ln( $n$ ); if ( $n \leq 0$ )  $\vee$  ( $n \geq \text{max\_n}$ ) then finish;
  for  $i \leftarrow 1$  to  $m$  do
    begin for  $j \leftarrow 1$  to  $n$  do
      begin read( $c$ );  $v[i,j] \leftarrow dd[c]$ ;
      end;
     $v[i,0] \leftarrow v[i,1]$ ;  $v[i,n+1] \leftarrow v[i,n]$ ;
    read_ln;
    end;
  for  $j \leftarrow 0$  to  $n+1$  do
    begin  $v[0,j] \leftarrow v[1,j]$ ;  $v[m+1,j] \leftarrow v[m,j]$ ;
    end

```

This code is used in section 23.

13. The code just written makes use of three temporary registers that must be declared:

```

⟨Global variables  $\epsilon$ ⟩ +≡
 $i, j$ : integer; { current row and column }
 $c$ : char; { character read from input }

```

14. **Pixel compensation.** The 33-level output of this program is assumed to be printed by a font that contains 4×8 characters, where each character has 0 to 32 black bits. Physical properties of output devices cause distortions, so that a character with k black bits does not have an apparent density of $k/32$. We therefore maintain a table of apparent density values.

```

define  $\text{max\_l} = 32$  { maximum output level }

```

```

⟨Global variables  $\epsilon$ ⟩ +≡
 $d$ : array [0 ..  $\text{max\_l}$ ] of real; { apparent densities, from 0.0 to 1.0 }

```

15. This table is based on some densitometer measurements that are not especially reliable. The amount of toner seems to vary between the top of a page and the bottom; also blocks of the character "N" seem to appear darker than blocks of the character "O", because of some property of xerography, although the "O" has one more bit turned on. Such anomalies have been smoothed out here, since the resulting values should prove good enough in practice.

```

⟨Set initial values  $\tau$ ⟩ +≡
 $d[0] \leftarrow 0.0$ ;  $d[1] \leftarrow 0.06$ ;  $d[2] \leftarrow 0.095$ ;  $d[3] \leftarrow 0.125$ ;  $d[4] \leftarrow 0.155$ ;
 $d[5] \leftarrow 0.175$ ;  $d[6] \leftarrow 0.215$ ;  $d[7] \leftarrow 0.245$ ;  $d[8] \leftarrow 0.27$ ;  $d[9] \leftarrow 0.29$ ;
 $d[10] \leftarrow 0.3$ ;  $d[11] \leftarrow 0.31$ ;  $d[12] \leftarrow 0.32$ ;  $d[13] \leftarrow 0.33$ ;  $d[14] \leftarrow 0.34$ ;
 $d[15] \leftarrow 0.35$ ;  $d[16] \leftarrow 0.36$ ;  $d[17] \leftarrow 0.37$ ;  $d[18] \leftarrow 0.38$ ;  $d[19] \leftarrow 0.4$ ;
 $d[20] \leftarrow 0.42$ ;  $d[21] \leftarrow 0.44$ ;  $d[22] \leftarrow 0.47$ ;  $d[23] \leftarrow 0.5$ ;  $d[24] \leftarrow 0.53$ ;
 $d[25] \leftarrow 0.57$ ;  $d[26] \leftarrow 0.61$ ;  $d[27] \leftarrow 0.66$ ;  $d[28] \leftarrow 0.72$ ;  $d[29] \leftarrow 0.80$ ;
 $d[30] \leftarrow 0.88$ ;  $d[31] \leftarrow 0.96$ ;  $d[32] \leftarrow 1.0$ ;

```


16. We convert the pixel values by using a variant of the Floyd-Steinberg algorithm for adaptive grayscale [Society for Information Display, *SID 75 Digest*, 36–37]. The idea is to find the best available density, then to diffuse the error into adjacent pixels that haven't yet been processed.

The following code assumes that x is the desired density value in column j of the current halfline. It outputs one 33-level density, then updates x and j in preparation for the next column. Adjustments to the densities in the two next halflines are accumulated in auxiliary arrays $next1$ and $next2$; this will compensate for errors in the current halfline.

```
We assume that  $next1[j]$ ,  $next1[j+1]$ , and  $next2[j]$  correspond to the dots that are adjacent to  $current[j]$ .
<Output one value and move to the next column 16> ≡
  <Find  $l$  so that  $d[l]$  is as close as possible to  $x$  21>;
  write(xchr["0" + l]); err ←  $x - d[l]$ ;
  next1[j] ← next1[j] + alpha * err;
  next2[j] ← beta * err;
  j ← j + 1; { move right }
  next1[j] ← next1[j] + gamma * err;
  x ← current[j] + delta * err
```

This code is used in sections 18 and 18.

17. The constants $alpha$.. $delta$ control the distribution of errors to adjacent dot positions.

```
<Set initial values 7> +≡
  alpha ← 7/16; { error diffusion to SW neighbor }
  beta ← 1/16; { error diffusion to S neighbor }
  gamma ← 5/16; { error diffusion to SE neighbor }
  delta ← 3/16; { error diffusion to E neighbor }
```

18. Here is the overall control of the process. Every halfline of the picture being output is a sequence of ASCII characters from "0" to "P", terminated by ".".

```
<Output the picture 18> ≡
  for j ← 1 to n + 1 do
    begin next1[j] ← 0.0; next2[j] ← 0.0;
    end;
  for i ← 1 to m do
    begin <Set the current halfline data for the upper row of dots in line i 19>;
    j ← 1; x ← current[1];
    repeat <Output one value and move to the next column 16>;
    until j > n;
    write_ln(" . "); <Set the current halfline data for the lower row of dots in line i 20>;
    j ← 1; x ← current[1];
    repeat <Output one value and move to the next column 16>;
    until j > n;
    write_ln(" . ");
    end
```

This code is used in section 23.

19. The density value for dot j in the upper halfline of line i is obtained as a weighted average of the input values in rows $i - 1$ and i , columns j and $j + 1$. The upper halfline is skewed to the right, so we must shift $next1$ and $next2$ appropriately.

```
<Set the current halfline data for the upper row of dots in line i 19> ≡
  for j ← 1 to n do
    begin current[j] ← (9 * v[i, j] + 3 * v[i, j + 1] + 3 * v[i - 1, j] + v[i - 1, j + 1])/16 + next1[j + 1];
    next1[j] ← next2[j];
    end;
  next1[n + 1] ← 0.0
```

This code is used in section 18.

20. The lower halfline is similar, but in this case there is leftward skew; we use rows i and $i + 1$, columns $j - 1$ and j .

```

⟨Set the current halfline data for the lower row of dots in line  $i$  20⟩ ≡
  for  $j \leftarrow 1$  to  $n$  do
    begin  $current[j] \leftarrow (9 * v[i, j] + 3 * v[i, j - 1] + 3 * v[i + 1, j] + v[i + 1, j - 1]) / 16 + next1[j]$ ;
       $next1[j + 1] \leftarrow next2[j]$ ;
    end;
   $next1[1] \leftarrow 0.0$ 

```

This code is used in section 18.

21. The algorithm is now complete except for the part that chooses the closest possible dot size. A straightforward binary search works well for this purpose:

```

⟨Find  $l$  so that  $d[l]$  is as close as possible to  $x$  21⟩ ≡
  if  $x \leq 0.0$  then  $l \leftarrow 0$ 
  else if  $x \geq 1.0$  then  $l \leftarrow max\_l$ 
  else begin  $low\_l \leftarrow 0$ ;  $high\_l \leftarrow max\_l$ ; { we have  $d[low\_l] \leq x < d[high\_l]$  }
    while  $high\_l - low\_l > 1$  do
      begin  $mid\_l \leftarrow (low\_l + high\_l) \text{ div } 2$ ;
        if  $x \geq d[mid\_l]$  then  $low\_l \leftarrow mid\_l$ 
        else  $high\_l \leftarrow mid\_l$ ;
      end;
    if  $x - d[low\_l] \leq d[high\_l] - x$  then  $l \leftarrow low\_l$  else  $l \leftarrow high\_l$ ;
  end

```

This code is used in section 16.

22. We had better declare the variables we've been using.

```

⟨Global variables 6⟩ +≡
 $x$ : real; { current pixel density }
 $err$ : real; { difference between  $x$  and the best we can achieve }
 $current$ : array [0 ..  $max\_n$ ] of real; { desired densities in current halfline }
 $next1, next2$ : array [0 ..  $max\_n$ ] of real; { corrections to subsequent densities }
 $alpha, beta, gamma, delta$ : real; { constants of error diffusion }
 $l, low\_l, mid\_l, high\_l$ : 0 ..  $max\_l$ ; { trial density levels }

```

23. The main program. Now we're ready to put all the pieces together.

```

⟨The main program 23⟩ ≡
  write_ln(`\input_hf33`); write_ln;
  while true do
    begin ⟨Input a picture, or terminate the program 12⟩;
      write_ln(`\beginhalftone`); ⟨Output the picture 18⟩;
      write_ln(`\endhalftone`); write_ln;
    end;
  cleanup_and_terminate;

```

This code is used in section 2.

Appendix 4: Pixel optimization

Here is another short WEB program. It was used to generate the special font for Mona Lisa.

1. Introduction. This program prepares a METAFONT program for a special-purpose font that will approximate a given picture. The input is assumed to be a binary file that contains one byte of density information per pixel. The output will be a sequence of lines like

```
row(10); cols(3,15,16,17);
```

this means that bits 3, 15, 16, and 17 of the character for row 10 should be black.

2. Here's an outline of the entire Pascal program:

```
program picfont(bytes_in, output);
  type <Types in the outer block 5>
  var <Global variables 6>
      <Basic procedures 10>
  begin <The main program 26>;
  end.
```

3. The picture in the input data is assumed to contain mm rows and nn columns.

```
define mm = 512 { this many rows }
define nn = 440 { this many columns }
```

4. It's convenient to declare a macro for incrementation.

```
define incr(#) ≡ # ← # + 1
```

5. Inputting and outputting the data. The input appears in a file of 8-bit bytes, with 00 representing black and FF representing white. There are $mm \times nn$ bytes; they appear in order from top to bottom and left to right just as we normally read a page of text.

```
<Types in the outer block 5> ≡
  eight_bits = 0 .. 255; { unsigned one-byte quantity }
  byte_file = packed file of eight_bits; { files that contain binary data }
```

This code is used in section 2.

```
6. <Global variables 6> ≡
bytes_in: byte_file;
```

See also sections 9, 14, 16, 22, and 25.

This code is used in section 2.

7. Different Pascal systems have different ways of dealing with binary files. Here is one common way.

```
<Open the input file 7> ≡
  reset(bytes_in, '^', '/B:8')
```

This code is used in section 26.

8. We shall use the following model for estimating the effect of a given bit pattern: If a pixel is black, the darkness is 1.0; if it is white but at least one of its four neighbors is black, the darkness is $zeta$; if it is white and has four white neighbors, the darkness is zero.

```
define white = 0 { code for a white pixel with all white neighbors }
define gray = 1 { code for a white pixel with 1, 2, 3, or 4 black neighbors }
define black = 2 { code for a black pixel }
define zeta ≡ 0.2 { assumed darkness of white pixel with a black neighbor }
```

9. There isn't room to store all the input bytes in memory at once, but it suffices to keep buffers for about a dozen rows near the current area being computed.

```

⟨Global variables 6⟩ +≡
ii: integer; { the buffer holds rows  $8ii - 7$  through  $8ii + 4$  }
buffer: array [-2 .. 9, 0 .. nn + 1] of real; { densities in twelve current rows }
darkness: array [-3 .. 9, 0 .. nn + 1] of white .. black; { darknesses in buffer rows }
new_row: array [0 .. nn + 1] of real; { densities in row being input }

```

10. The *get_in* procedure computes the densities in a specified row and puts them in *new_row*. This procedure is called successively for $i = 1, 2, \dots$

```

⟨Basic procedures 10⟩ ≡
procedure get_in(i: integer);
  var j: integer; t: eight_bits; { byte of input }
  begin new_row[0] ← 0.0;
  if i > mm then
    for j ← 1 to nn do new_row[j] ← 0.0
  else for j ← 1 to nn do
    begin read(bytes_in, t); new_row[j] ← (255.5 - t)/256.0;
    end;
  new_row[nn + 1] ← 0.0;
  end;

```

See also sections 11 and 20.

This code is used in section 2.

11. Here is a procedure that "rolls" the buffer down eight lines:

```

⟨Basic procedures 10⟩ +≡
procedure roll;
  var j: 0 .. nn + 1; i: 2 .. 9; k: integer;
  begin for i ← 6 to 9 do
    for j ← 0 to nn + 1 do
      begin buffer[i - 8, j] ← buffer[i, j]; darkness[i - 8, j] ← darkness[i, j];
      end;
    for j ← 0 to nn + 1 do darkness[-3, j] ← darkness[5, j];
    incr(ii);
  for i ← 2 to 9 do
    begin get_in( $8 * ii + i - 5$ );
    for j ← 0 to nn + 1 do
      begin buffer[i, j] ← new_row[j]; darkness[i, j] ← white;
      end;
    end;
  end;

```

12. It's tedious but not difficult to get everything started. We put zeros above the top lines in the picture.

```

⟨Initialize the buffers 12⟩ ≡
  ii ← 0;
  for i ← 6 to 9 do
    begin get_in(i - 5);
    for j ← 0 to nn + 1 do
      begin buffer[i, j] ← new_row[j]; darkness[i, j] ← white;
      end;
    end;
  for i ← -2 to 5 do
    for j ← 0 to nn + 1 do
      begin buffer[i, j] ← 0.0; darkness[i, j] ← white;
      end;
    for j ← 0 to nn + 1 do darkness[-3, j] ← white

```

This code is used in section 26.

13. It's easy to output the current darkness values. Here we output eight consecutive rows.

⟨Output the pixel values for the top eight rows of the buffer 13⟩ ≡

```

for i ← -2 to 5 do
  begin write('row(', 8 * ii - 5 + i : 1, '); cols('); cols_out ← 0;
  for j ← 1 to nn do
    if darkness[i, j] = black then
      begin if cols_out < 15 then
        begin if cols_out > 0 then write(', ');
          incr(cols_out);
        end
      else begin write_ln(', '); write(' '); cols_out ← 1;
        end;
      write(j : 1);
    end;
  write_ln(');')
end

```

This code is used in section 26.

14. ⟨Global variables 6⟩ +≡

cols_out: 0 .. 15; { the number of columns output so far on this line }

15. **Dot diffusion.** The pixels are divided into 64 classes, numbered from 0 to 63. We convert the pixel values to darknesses by using a method called "dot diffusion." Values are assigned first to all the pixels of class 0, then to all the pixels of class 1, etc.; the error incurred at each step is distributed to the neighbors whose class numbers are higher. This is done by means of precomputed tables *class_row*, *class_col*, *start*, *del_i*, *del_j*, and *alpha* whose function is easy to deduce from the following code:

⟨Choose pixel values and diffuse the errors in the buffer 15⟩ ≡

```

for k ← 0 to 63 do
  begin i ← class_row[k]; j ← class_col[k];
  while j ≤ nn do
    begin ⟨Decide the color of pixel [i, j] and the resulting err 17);
    for l ← start[k] to start[k + 1] - 1 do
      begin u ← i + del_i[l]; v ← j + del_j[l]; buffer[u, v] ← buffer[u, v] + err * alpha[l];
      end;
    j ← j + 8;
  end;
end

```

This code is used in section 26.

16. ⟨Global variables 6⟩ +≡

class_row: **array** [0 .. 63] **of** -2 .. 8; { buffer row containing pixels of a given class }

class_col: **array** [0 .. 63] **of** 1 .. 8; { first column containing pixels of a given class }

class_number: **array** [-2 .. 9, 0 .. 9] **of** 0 .. 63; { number of a given position }

err: *real*; { error introduced at the current position }

err_black: *real*; { error introduced at the current position if black chosen }

black_diff: *real*; { difference between *err* and *err_black* for gray pixel }

l: 0 .. 256; { index into diffusion tables }

start: **array** [0 .. 64] **of** 0 .. 256; { first entry of diffusion table for a given class }

del_i, *del_j*: **array** [0 .. 256] **of** -1 .. 1; { neighboring location for diffusion }

alpha: **array** [0 .. 256] **of** *real*; { constant of proportionality for diffusion }

17. Here we choose white or black, whichever minimizes the magnitude of the error. Potentially *gray* values of this pixel and its neighbors make this calculation slightly tricky, as we must subtract *zeta* when a gray pixel is created and add *zeta* when it is destroyed.

```

⟨Decide the color of pixel [i,j] and the resulting err 17⟩ ≡
  if darkness[i,j] = gray then
    begin err ← buffer[i,j] - zeta; err_black ← err - black_diff;
    end
  else begin err ← buffer[i,j]; err_black ← err - 1.0;
  end;
  if darkness[i-1,j] = white then err_black ← err_black - zeta;
  if darkness[i,j-1] = white then err_black ← err_black - zeta;
  if darkness[i,j+1] = white then err_black ← err_black - zeta;
  if darkness[i+1,j] = white then err_black ← err_black - zeta;
  if err_black + err > 0 then
    begin err ← err_black; darkness[i,j] ← black;
    if darkness[i-1,j] = white then darkness[i-1,j] ← gray;
    if darkness[i,j-1] = white then darkness[i,j-1] ← gray;
    if darkness[i,j+1] = white then darkness[i,j+1] ← gray;
    if darkness[i+1,j] = white then darkness[i+1,j] ← gray;
    end

```

This code is used in section 15.

18. ⟨Initialize the diffusion tables 18⟩ ≡
black_diff ← 1.0 - 2.0 * *zeta*;

See also section 19.

This code is used in section 26.

19. **Computing the diffusion tables.** The tables for dot diffusion could be specified by a large number of boring assignment statements, but it is more fun to compute them by a method that shows some of the mysterious underlying structure.

```

⟨Initialize the diffusion tables 18⟩ +≡
  ⟨Initialize the class number matrix 21⟩;
  ⟨Compile "instructions" for the diffusion operations 23⟩

```

20. The order of classes used here is the order in which pixels might be blackened in a font for halftones based on dots in a 45° grid. In fact, this is precisely the pattern used in the —dot300— font that was described earlier.

```

⟨Basic procedures 10⟩ +≡
procedure store(i,j : integer); { establish new class_row, class_col }
  begin if i < 1 then i ← i + 8 else if i > 8 then i ← i - 8;
  if j < 1 then j ← j + 8 else if j > 8 then j ← j - 8;
  class_number[i,j] ← k; class_row[k] ← i; class_col[k] ← j; incr(k);
  end;
procedure store_eight(i,j : integer); { rotate and shift for eight classes }
  begin store(i,j); store(i-4,j+4); store(5-j,i); store(1-j,i-4);
  store(4+j,1-i); store(j,5-i); store(5-i,5-j); store(1-i,1-j);
  end;

```

21. \langle Initialize the class number matrix 21 $\rangle \equiv$
 $k \leftarrow 0$; $store_eight(7,2)$; $store_eight(8,3)$; $store_eight(8,2)$; $store_eight(8,1)$;
 $store_eight(1,4)$; $store_eight(1,3)$; $store_eight(1,2)$; $store_eight(2,3)$;
for $i \leftarrow 1$ **to** 8 **do**
 begin $class_number[i,0] \leftarrow class_number[i,8]$; $class_number[i,9] \leftarrow class_number[i,1]$;
 end;
for $j \leftarrow 0$ **to** 9 **do**
 begin $class_number[-2,j] \leftarrow class_number[6,j]$; $class_number[-1,j] \leftarrow class_number[7,j]$;
 $class_number[0,j] \leftarrow class_number[8,j]$; $class_number[9,j] \leftarrow class_number[1,j]$;
 end

This code is used in section 19.

22. The tricky part of this process is the fact that some values near the bottom of the buffer aren't ready for processing until errors have been diffused from the next bufferload. In such cases we go up eight rows to process a value that has been held over.

\langle Global variables 6 $\rangle + \equiv$
 $hold$: **array** [0 .. 9, 0 .. 9] **of** *boolean*;
 { is this value too close to the bottom of the buffer to allow immediate processing? }

23. The "compilation" in this step simulates going through the diffusion process the slow way, and records the actions it performs (so that they can all be done at high speed later).

\langle Compile "instructions" for the diffusion operations 23 $\rangle \equiv$
 for $j \leftarrow 0$ **to** 9 **do** $hold[9,j] \leftarrow true$;
 for $i \leftarrow 0$ **to** 8 **do**
 for $j \leftarrow 0$ **to** 9 **do** $hold[i,j] \leftarrow false$;
 $l \leftarrow 0$; $k \leftarrow 0$;
 repeat $i \leftarrow class_row[k]$; $j \leftarrow class_col[k]$; $w \leftarrow 0$; $start[k] \leftarrow l$;
 for $u \leftarrow i - 1$ **to** $i + 1$ **do**
 for $v \leftarrow j - 1$ **to** $j + 1$ **do**
 if $class_number[u,v] > k$ **then**
 begin $del_i[l] \leftarrow u - i$; $del_j[l] \leftarrow v - j$; $incr(l)$;
 if $u = i$ **then** $w \leftarrow w + 2$ { neighbors in the same row get weight 2 }
 else if $v = j$ **then** $w \leftarrow w + 2$ { neighbors in the same column get weight 2 }
 else $incr(w)$; { diagonal neighbors get weight 1 }
 end
 else if $hold[u,v]$ **then** $hold[i,j] \leftarrow true$;
 if $hold[i,j]$ **then** $class_row[k] \leftarrow i - 8$;
 \langle Compute the *alpha* values for class k , given the total weight w 24 \rangle ;
 $incr(k)$;
 until $k = 64$;
 $start[64] \leftarrow l$

This code is used in section 19.

24. \langle Compute the *alpha* values for class k , given the total weight w 24 $\rangle \equiv$
 for $ll \leftarrow start[k]$ **to** $l - 1$ **do**
 begin if $del_i[ll] = 0$ **then** $alpha[ll] \leftarrow 2.0/w$
 else if $del_j[ll] = 0$ **then** $alpha[ll] \leftarrow 2.0/w$
 else $alpha[ll] \leftarrow 1.0/w$;
 end

This code is used in section 23.

25. \langle Global variables 6 $\rangle + \equiv$
 ll : 0 .. 256; { loop index }
 u, v : *integer*; { neighbors of i and j }
 w : *integer*; { the weighted number of high-class neighbors }
 i, j : *integer*; { the current pixel position being considered }
 k : 0 .. 64; { the current class being considered }

26. The main program. Finally we're ready to get it all together.


```

<The main program 26> ≡
  <Initialize the diffusion tables 18>;
  <Open the input file 7>;
  <Initialize the buffers 12>;
  repeat <Choose pixel values and diffuse the errors in the buffer 15>;
    if ii > 0 then <Output the pixel values for the top eight rows of the buffer 13>;
    roll;
  until 8 * ii > mm

```

This code is used in section 2.

Acknowledgements. The research described in this paper was supported in part by the System Development Foundation and in part by National Science Foundation grants IST-8201926, MCS-8300984, and DCR-8308109.



Addendum

Stop the presses! When I wrote the preceding pages (and had them typeset), I was unaware of a “well known” method that should have been included for comparison. So far this paper has considered (1) a halftone font with 65 levels of gray, in which each 8×8 character essentially contributes two dots to a picture; and (2) a halftone font with 33 levels of gray, in which each 4×8 character contributes one dot to a picture. It’s also possible to construct (3) a halftone font with 17 levels of gray, in which each 4×4 character essentially contributes half of a dot (actually two quarter-dots) to a picture. This third method is based on an idea due to Robert L. Gard [*Computer Graphics and Image Processing* 5 (1976), 151–171].

The k th level of gray in the half-dot scheme is obtained by blackening cells 0 to $k - 1$ in the array

1	5	10	14	or	14	10	5	1
3	7	8	12		12	8	7	3
13	9	6	2		2	6	9	13
15	11	4	0		0	4	11	15

(We actually make two sets of characters, one the mirror image of the other, and alternate between them as a picture is typeset.) The following METAFONT file will generate such a font `hd300`, in essentially the same way that the other fonts `dot300` and `hf300` were generated earlier:

```
% halftone font with 17 levels of gray, characters "A" (white) to "Q" (black)
% includes also the mirror-reflected characters "a" (white) to "q" (black)
pair p[]; % the pixels in order (first p0 becomes black, then p1, etc.)
p0=(3,0); p4=(2,0); p8=(2,2); p12=(3,2);
transform r; r=identity rotatedaround ((1.5,1.5),180);
for i=0 step 4 until 12: p[i+1]=p[i] transformed r;
  p[i+2]=p[i] shifted (0,1); p[i+3]=p[i+2] transformed r; endfor
w#:=4/pt; % that's 4 pixels
font_quad:=w#; designsize:=8w#;
r:=identity reflectedabout ((2,0),(2,3));
picture prevchar; prevchar=nullpicture; % the pixels blackened so far
for i=0 upto 16:
  beginchar(i+ASCII"A",w#,w#,0); currentpicture:=prevchar;
  if i>0: addto currentpicture also unitpixel shifted p[i-1]; fi
  prevchar:=currentpicture; endchar;
  beginchar(i+ASCII"a",w#,w#,0);
  currentpicture:=prevchar transformed r; endchar;
endfor
```

Here are four pictures for comparison, showing also the result of the elaborate “dot diffusion” method discussed at the end of my paper:



double dot



single dot



half dot



dot diffusion

