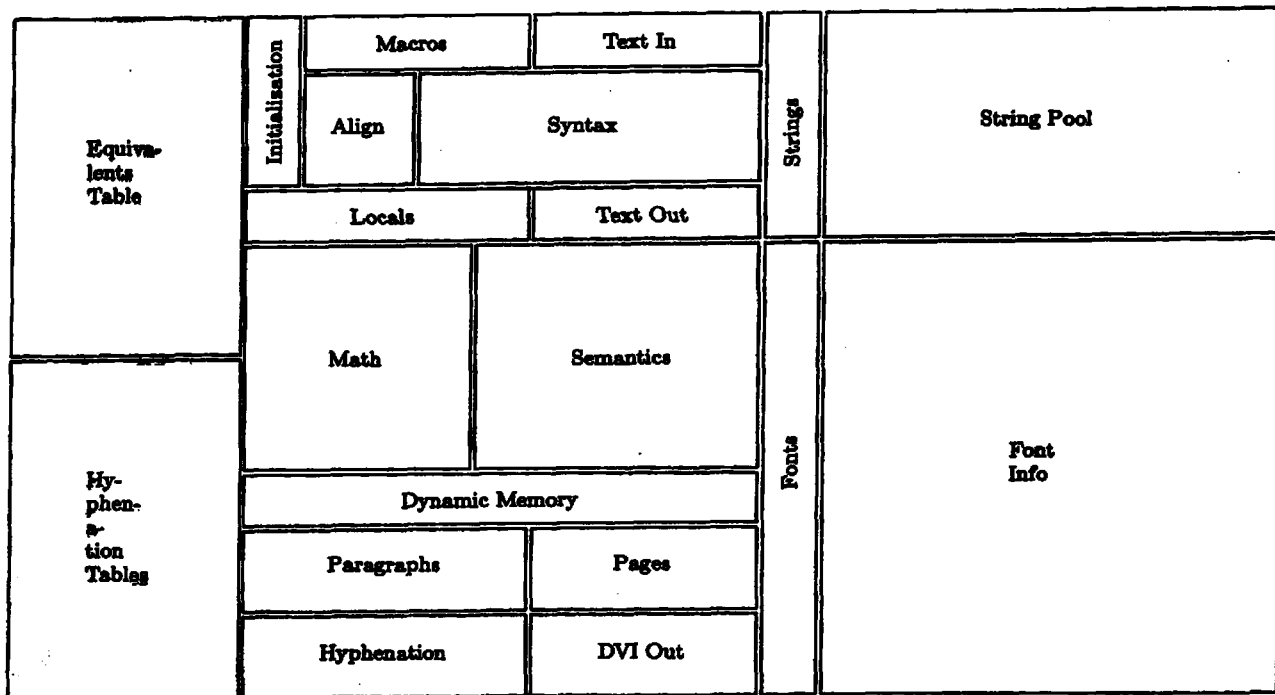


\* \* \* \* \*  
 Software  
 \* \* \* \* \*

**T<sub>E</sub>X<sub>82</sub> MEMORY STRUCTURE**

*Editor's note: The following diagram shows the relative amount of memory required by T<sub>E</sub>X<sub>82</sub> for various tables and work areas. It has been redrawn (not using T<sub>E</sub>X, although that would be a good problem) from a slide that Don Knuth prepared for the lecture on data structures at the T<sub>E</sub>X<sub>82</sub> Short Course held in conjunction with the July TUG meeting.*



□ ≈ 1K bytes

Memory represented in this diagram is approximately 260K bytes, plus (not shown) 100K of dynamic memory. On a DEC 20, T<sub>E</sub>X<sub>82</sub> occupies total memory of between 350K and 600K bytes, excluding Pascal run-time requirements.

\* \* \* \* \*

*Editor's note: David Fuchs has provided the following copy describing the format of .DVI files, including some changes made since the July TUG meeting. A summary of the changes appears in David's "News from Stanford" article, which begins on page 20.*

**1. Device-independent file format.** The most important output produced by a run of T<sub>E</sub>X is the “device independent” (DVI) file that specifies where characters and rules are to appear on printed pages. The form of these files was designed by David R. Fuchs in 1979. Almost any reasonable device can be driven by a program that takes DVI files as input, and dozens of such DVI-to-whatever programs have been written. Thus, it is possible to print the output of T<sub>E</sub>X on many different kinds of equipment, using T<sub>E</sub>X as a device-independent “front end.”

A DVI file is a stream of 8-bit bytes, which may be regarded as a series of commands in a machine-like language. The first byte of each command is the operation code, and this code is followed by zero or more bytes that provide parameters to the command. The parameters themselves may consist of several consecutive bytes; for example, the ‘*set\_rule*’ command has two parameters, each of which is four bytes long. Parameters are usually regarded as nonnegative integers; but four-byte-long parameters, and shorter parameters that denote distances, can be either positive or negative. Such parameters are given in two’s complement notation. For example, a two-byte-long distance parameter has a value between  $-2^{15}$  and  $2^{15} - 1$ .

A DVI file consists of a “preamble,” followed by a sequence of one or more “pages,” followed by a “postamble.” The preamble is simply a *pre* command, with its parameters that define the dimensions used in the file; this must come first. Each “page” consists of a *bop* command, followed by any number of other commands that tell where characters are to be placed on a physical page, followed by an *epc* command. The pages appear in the order that T<sub>E</sub>X generated them. If we ignore *nop* commands and *fmt\_def* commands (which are allowed between any two commands in the file), each *epc* command is immediately followed by a *bop* command, or by a *post* command; in the latter case, there are no more pages in the file, and the remaining bytes form the postamble. Further details about the postamble will be explained later.

Some parameters in DVI commands are “pointers.” These are four-byte quantities that give the location number of some other byte in the file; the first byte is number 0, then comes number 1, and so on. For example, one of the parameters of a *bop* command points to the previous *bop*; this makes it feasible to read the pages in backwards order, in case the results are being directed to a device that stacks its output face up. Suppose the preamble of a DVI file occupies bytes 0 to 99. Now if the first page occupies bytes 100 to 999, say, and if the second page occupies bytes 1000 to 1999, then the *bop* that starts in byte 1000 points to 100 and the *bop* that starts in byte 2000 points to 1000. (The very first *bop*, i.e., the one that starts in byte 100, has a pointer of  $-1$ .)

**2.** The DVI format is intended to be both compact and easily interpreted by a machine. Compactness is achieved by making most of the information implicit instead of explicit; when a DVI-reading program reads the commands for a page, it keeps track of several quantities: (a) The current font *f* is an integer; this value is changed only by *fmt* and *fmt\_num* commands. (b) The current position on the page is given by two numbers called the horizontal and vertical coordinates, *h* and *v*. Both coordinates are zero at the upper left corner of the page; moving to the right corresponds to increasing the horizontal coordinate, and moving down corresponds to increasing the vertical coordinate. Thus, the coordinates are essentially Cartesian, except that vertical directions are flipped; the Cartesian version of (*h*, *v*) would be (*h*,  $-v$ ). (c) The current spacing amounts are given by four numbers *w*, *x*, *y*, and *z*, where *w* and *x* are used for horizontal spacing and where *y* and *z* are used for vertical spacing. (d) There is a stack containing (*h*, *v*, *w*, *x*, *y*, *z*) values; the DVI commands *push* and *pop* are used to change the current level of operation. Note that the current font *f* is not pushed and popped; the stack contains only information about positioning.

The values of *h*, *v*, *w*, *x*, *y*, and *z* are signed integers having up to 32 bits; including the sign. Since they represent physical distances, there is a small unit of measurement such that increasing *h* by 1 means moving a certain tiny distance to the right. The actual unit of measurement is variable, as explained below; T<sub>E</sub>X sets things up so that its DVI output is in sp units, i.e., scaled points, in agreement with all the *scaled* dimensions in T<sub>E</sub>X’s data structures.

3. Here is list of all the commands that may appear in a DVI file. With each command we give its symbolic name (e.g., *bop*), its opcode byte (e.g., 129), and its parameters (if any). The parameters are followed by a bracketed number telling how many bytes they occupy; for example, '*p*[4]' means that parameter *p* is four bytes long.

- set\_char\_0* 0. Typeset character number 0 from font *f* such that the reference point of the character is at (*h*, *v*). Then increase *h* by the width of that character. Note that a character may have zero or negative width, so one cannot be sure that *h* will advance after this command; but *h* usually does increase.
- set\_char\_1* through *set\_char\_127* (opcodes 1 to 127). Do the operations of *set\_char\_0*, but use the appropriate character number instead of character 0.
- set1* 128 *c*[1]. Same as *set\_char\_0*, except that character number *c* is typeset. T<sub>E</sub>X82 uses this command for characters in the range  $128 \leq c < 256$ .
- set2* 129 *c*[2]. Same as *set1*, except that *c* is two bytes long, so it is in the range  $0 \leq c < 65536$ . T<sub>E</sub>X82 never uses this command, but it should come in handy for extensions of T<sub>E</sub>X that deal with oriental languages.
- set3* 130 *c*[3]. Same as *set1*, except that *c* is three bytes long, so it can be as large as  $2^{24} - 1$ . Not even the Chinese language has this many characters, but this command might prove useful in some yet unforeseen extension.
- set4* 131 *c*[4]. Same as *set1*, except that *c* is four bytes long. Imagine that.
- set\_rule* 132 *a*[4] *b*[4]. Typeset a solid black rectangle of height *a* and width *b*, with its bottom left corner at (*h*, *v*). Then set  $h \leftarrow h + b$ . If either  $a \leq 0$  or  $b \leq 0$ , nothing should be typeset. Note that if  $b < 0$ , the value of *h* will decrease even though nothing else happens. See below for details about how to typeset rules so that consistency with METAFONT is guaranteed.
- put1* 133 *c*[1]. Typeset character number *c* from font *f* such that the reference point of the character is at (*h*, *v*). (The 'put' commands are exactly like the 'set' commands, except that they simply put out a character or a rule without moving the reference point afterwards.)
- put2* 134 *c*[2]. Same as *set2*, except that *h* is not changed.
- put3* 135 *c*[3]. Same as *set3*, except that *h* is not changed.
- put4* 136 *c*[4]. Same as *set4*, except that *h* is not changed.
- put\_rule* 137 *a*[4] *b*[4]. Same as *set\_rule*, except that *h* is not changed.
- nop* 138. No operation, do nothing. Any number of *nop*'s may occur between DVI commands, but a *nop* cannot be inserted between a command and its parameters or between two parameters.
- bop* 139 *c*<sub>0</sub>[4] *c*<sub>1</sub>[4] ... *c*<sub>9</sub>[4] *p*[4]. Beginning of a page: Set  $(h, v, w, x, y, z) \leftarrow (0, 0, 0, 0, 0, 0)$  and set the stack empty. Set the current font *f* to an undefined value. The ten *c*<sub>*i*</sub> parameters hold the values of `\count0` ... `\count9` in T<sub>E</sub>X at the time `\shipout` was invoked for this page; they can be used to identify pages, if a user wants to print only part of a DVI file. The parameter *p* points to the previous *bop* command in the file, where the first *bop* has  $p = -1$ .
- eop* 140. End of page: Print what you have read since the previous *bop*. At this point the stack should be empty. (The DVI-reading programs that drive most output devices will have kept a buffer of the material that appears on the page that has just ended. This material is largely, but not entirely, in order by *v* coordinate and (for fixed *v*) by *h* coordinate; so it usually needs to be sorted into some order that is appropriate for the device in question.)
- push* 141. Push the current values of  $(h, v, w, x, y, z)$  onto the top of the stack; do not change any of these values. Note that *f* is not pushed.
- pop* 142. Pop the top six values off of the stack and assign them respectively to  $(h, v, w, x, y, z)$ . The number of pops should never exceed the number of pushes, since it would be highly embarrassing if the stack were empty at the time of a *pop* command.
- right1* 143 *b*[1]. Set  $h \leftarrow h + b$ , i.e., move right *b* units. The parameter is a signed number in two's complement notation,  $-128 \leq b < 128$ ; if  $b < 0$ , the reference point actually moves left.

## 4 DEVICE-INDEPENDENT FILE FORMAT

T<sub>E</sub>X82 3053

- right2* 144 *b*[2]. Same as *right1*, except that *b* is a two-byte quantity in the range  $-32768 \leq b < 32768$ .
- right3* 145 *b*[3]. Same as *right1*, except that *b* is a three-byte quantity in the range  $-2^{23} \leq b < 2^{23}$ .
- right4* 146 *b*[4]. Same as *right1*, except that *b* is a four-byte quantity in the range  $-2^{31} \leq b < 2^{31}$ .
- w0* 147. Set  $h \leftarrow h + w$ ; i.e., move right *w* units. With luck, this parameterless command will usually suffice, because the same kind of motion will occur several times in succession; the following commands explain how *w* gets particular values.
- w1* 148 *b*[1]. Set  $w \leftarrow b$  and  $h \leftarrow h + b$ . The value of *b* is a signed quantity in two's complement notation,  $-128 \leq b < 128$ . This command changes the current *w* spacing and moves right by *b*.
- w2* 149 *b*[2]. Same as *w1*, but *b* is two bytes long,  $-32768 \leq b < 32768$ .
- w3* 150 *b*[3]. Same as *w1*, but *b* is three bytes long,  $-2^{23} \leq b < 2^{23}$ .
- w4* 151 *b*[4]. Same as *w1*, but *b* is four bytes long,  $-2^{31} \leq b < 2^{31}$ .
- x0* 152. Set  $h \leftarrow h + x$ ; i.e., move right *x* units. The '*x*' commands are like the '*w*' commands except that they involve *x* instead of *w*.
- x1* 153 *b*[1]. Set  $x \leftarrow b$  and  $h \leftarrow h + b$ . The value of *b* is a signed quantity in two's complement notation,  $-128 \leq b < 128$ . This command changes the current *x* spacing and moves right by *b*.
- x2* 154 *b*[2]. Same as *x1*, but *b* is two bytes long,  $-32768 \leq b < 32768$ .
- x3* 155 *b*[3]. Same as *x1*, but *b* is three bytes long,  $-2^{23} \leq b < 2^{23}$ .
- x4* 156 *b*[4]. Same as *x1*, but *b* is four bytes long,  $-2^{31} \leq b < 2^{31}$ .
- down1* 157 *a*[1]. Set  $v \leftarrow v + a$ , i.e., move down *a* units. The parameter is a signed number in two's complement notation,  $-128 \leq a < 128$ ; if *a* < 0, the reference point actually moves up.
- down2* 158 *a*[2]. Same as *down1*, except that *a* is a two-byte quantity in the range  $-32768 \leq a < 32768$ .
- down3* 159 *a*[3]. Same as *down1*, except that *a* is a three-byte quantity in the range  $-2^{23} \leq a < 2^{23}$ .
- down4* 160 *a*[4]. Same as *down1*, except that *a* is a four-byte quantity in the range  $-2^{31} \leq a < 2^{31}$ .
- y0* 161. Set  $v \leftarrow v + y$ ; i.e., move down *y* units. With luck, this parameterless command will usually suffice, because the same kind of motion will occur several times in succession; the following commands explain how *y* gets particular values.
- y1* 162 *a*[1]. Set  $y \leftarrow a$  and  $v \leftarrow v + a$ . The value of *a* is a signed quantity in two's complement notation,  $-128 \leq a < 128$ . This command changes the current *y* spacing and moves down by *a*.
- y2* 163 *a*[2]. Same as *y1*, but *a* is two bytes long,  $-32768 \leq a < 32768$ .
- y3* 164 *a*[3]. Same as *y1*, but *a* is three bytes long,  $-2^{23} \leq a < 2^{23}$ .
- y4* 165 *a*[4]. Same as *y1*, but *a* is four bytes long,  $-2^{31} \leq a < 2^{31}$ .
- z0* 166. Set  $v \leftarrow v + z$ ; i.e., move down *z* units. The '*z*' commands are like the '*y*' commands except that they involve *z* instead of *y*.
- z1* 167 *a*[1]. Set  $z \leftarrow a$  and  $v \leftarrow v + a$ . The value of *a* is a signed quantity in two's complement notation,  $-128 \leq a < 128$ . This command changes the current *z* spacing and moves down by *a*.
- z2* 168 *a*[2]. Same as *z1*, but *a* is two bytes long,  $-32768 \leq a < 32768$ .
- z3* 169 *a*[3]. Same as *z1*, but *a* is three bytes long,  $-2^{23} \leq a < 2^{23}$ .
- z4* 170 *a*[4]. Same as *z1*, but *a* is four bytes long,  $-2^{31} \leq a < 2^{31}$ .
- font\_num\_0* 171. Set  $f \leftarrow 0$ . Font 0 must previously have been defined by a *font\_def* instruction, as explained below.
- font\_num\_1* through *font\_num\_63* (opcodes 172 to 234). Set  $f \leftarrow 1, \dots, f \leftarrow 63$ , respectively.
- font1* 235 *k*[1]. Set  $f \leftarrow k$ . T<sub>E</sub>X82 uses this command for font numbers in the range  $64 \leq k < 256$ .
- font2* 236 *k*[2]. Same as *font1*, except that *k* is two bytes long, so it is in the range  $0 \leq k < 65536$ . T<sub>E</sub>X82 never generates this command, but large font numbers may prove useful for specifications of color or texture, or they may be used for special fonts that have fixed numbers in some external coding scheme.

- font* 237 *k*[3]. Same as *font*, except that *k* is three bytes long, so it can be as large as  $2^{24} - 1$ .
- font* 238 *k*[4]. Same as *font*, except that *k* is four bytes long; this is for the really big font numbers (and for the negative ones).
- xxx* 239 *k*[1] *x*[*k*]. This command is undefined in general; it functions as a (*k*+2)-byte *nop* unless special DVI-reading programs are being used. T<sub>E</sub>X82 generates *xxx* when a normal `\xsend` appears, setting *k* to the number of bytes being sent. It is recommended that *x* be a string having the form of a keyword followed by possible parameters relevant to that keyword.
- xxx* 240 *k*[2] *x*[*k*]. Like *xxx*, but  $0 \leq k < 65536$ .
- xxx* 241 *k*[3] *x*[*k*]. Like *xxx*, but  $0 \leq k < 2^{24}$ .
- xxx* 242 *k*[4] *x*[*k*]. Like *xxx*, but *k* can be ridiculously large. T<sub>E</sub>X82 uses *xxx* when sending a string of length 256 or more.
- font\_def* 243 *k*[1] *c*[4] *s*[4] *d*[4] *a*[1] *l*[1] *n*[*a*+*l*]. Define font *k*, where  $0 \leq k < 63$ ; font definitions will be explained shortly.
- font\_def* 244 *k*[2] *c*[4] *s*[4] *d*[4] *a*[1] *l*[1] *n*[*a*+*l*]. Define font *k*, where  $0 \leq k < 65536$ .
- font\_def* 245 *k*[3] *c*[4] *s*[4] *d*[4] *a*[1] *l*[1] *n*[*a*+*l*]. Define font *k*, where  $0 \leq k < 2^{24}$ .
- font\_def* 246 *k*[4] *c*[4] *s*[4] *d*[4] *a*[1] *l*[1] *n*[*a*+*l*]. Define font *k*, where  $-2^{31} \leq k < 2^{30}$ .
- pre* 247 *i*[1] *num*[4] *den*[4] *mag*[4] *k*[1] *x*[*k*]. Beginning of the preamble; this must come at the very beginning of the file. Parameters *i*, *num*, *den*, *mag*, *k*, and *x* are explained below.
- post* 248. Beginning of the postamble, see below.
- post\_post* 249. Ending of the postamble, see below.
- Commands 250-255 are undefined at the present time.

4. The preamble contains basic information about the file as a whole. As stated above, there are six parameters:

$$i[1] \text{ num}[4] \text{ den}[4] \text{ mag}[4] \text{ k}[1] \text{ x}[k].$$

The *i* byte identifies DVI format; currently this byte is always set to 2. (Some day we will set *i* = 3, when DVI format makes another incompatible change—perhaps in 1992.)

The next two parameters, *num* and *den*, are positive integers that define the units of measurement; they are the numerator and denominator of a fraction by which all dimensions in the DVI file could be multiplied in order to get lengths in units of  $10^{-7}$  meters. Since there are 72.27 points per inch and 2.54 centimeters per inch, and since T<sub>E</sub>X82 works with scaled points where there are  $2^{16}$  sp in a point, T<sub>E</sub>X82 sets *num* = 25400000 and *den* =  $7227 \cdot 2^{16} = 473628672$ .

The *mag* parameter is what T<sub>E</sub>X calls `\mag`, i.e., 1000 times the desired magnification. The actual fraction by which dimensions are multiplied is therefore *mag* · *num* / 1000*den*. Note that if a T<sub>E</sub>X source document does not call for any ‘true’ dimensions, and if you change it only by specifying a different `\mag` setting, the DVI file that T<sub>E</sub>X creates will be completely unchanged except for the value of *mag* in the preamble and postamble. (Fancy DVI-reading programs allow users to override the *mag* setting when a DVI file is being printed.)

Finally, *k* and *x* allow the DVI writer to include a comment, which is not interpreted further. The length of comment *x* is *k*, where  $0 \leq k < 256$ .

define *id\_byte* = 2 { identifies the kind of DVI files described here }

5. Font definitions for a given font number  $k$  contain further parameters
$$c[4] s[4] d[4] a[1] l[1] n[a + l].$$

The four-byte value  $c$  is the check sum that T<sub>E</sub>X found in the TFM file for this font;  $c$  should match the check sum of the font found by programs that read this DVI file.

Parameter  $s$  contains a fixed-point scale factor that is applied to the character widths in font  $k$ ; font dimensions in TFM files and other font files are relative to this quantity, which is called the “at size” elsewhere in this documentation. The value of  $s$  is always positive and less than  $2^{27}$ . It is given in the same units as the other DVI dimensions, i.e., in sp when T<sub>E</sub>X82 has made the file. Parameter  $d$  is similar to  $s$ ; it is the “design size,” and it is given in DVI units that have not been corrected for the magnification  $mag$  found in the preamble. Thus, font  $k$  is to be used at  $mag \cdot s/1000d$  times its normal size.

The remaining part of a font definition gives the external name of the font, which is an ascii string of length  $a + l$ . The number  $a$  is the length of the “area” or directory, and  $l$  is the length of the font name itself; the standard local system font `arcn` is supposed to be used when  $a = 0$ . The  $n$  field contains the area in its first  $a$  bytes.

Font definitions must appear before the first use of a particular font number. Once font  $k$  is defined, it must not be defined again; however, we shall see below that font definitions appear in the postamble as well as in the pages, so in this sense each font number is defined exactly twice, if at all. Like `nop` commands, font definitions can appear before the first `bop`, or between an `eop` and a `bop`.

6. Sometimes it is desirable to make horizontal or vertical rules line up precisely with certain features in characters of a font. It is possible to guarantee the correct matching between DVI output and the characters generated by METAFONT by adhering to the following principles: (1) The METAFONT characters should be positioned so that a bottom edge or left edge that is supposed to line up with the bottom or left edge of a rule appears at the reference point, i.e., in row 0 and column 0 of the METAFONT raster. This ensures that the position of the rule will not be rounded differently when the pixel size is not a perfect multiple of the units of measurement in the DVI file. (2) A typeset rule of height  $a > 0$  and width  $b > 0$  should be equivalent to a METAFONT-generated character having black pixels in precisely those raster positions whose METAFONT coordinates satisfy  $0 \leq x < ab$  and  $0 \leq y < \alpha a$ , where  $\alpha$  is the number of pixels per DVI unit.

7. The last page in a DVI file is followed by ‘`post`’; this command introduces the postamble, which summarizes important facts that T<sub>E</sub>X has accumulated about the file, making it possible to print subsets of the data with reasonable efficiency. The postamble has the form

$$\begin{aligned} & \text{post } p[4] \text{ num}[4] \text{ den}[4] \text{ mag}[4] l[4] u[4] s[2] t[2] \\ & \text{(font definitions)} \\ & \text{post\_post } q[4] i[1] 223's[\geq 4] \end{aligned}$$

Here  $p$  is a pointer to the final `bop` in the file. The next three parameters,  $num$ ,  $den$ , and  $mag$ , are duplicates of the quantities that appeared in the preamble.

Parameters  $l$  and  $u$  give respectively the height-plus-depth of the tallest page and the width of the widest page, in the same units as other dimensions of the file. These numbers might be used by a DVI-reading program to position individual “pages” on large sheets of film or paper.

Parameter  $s$  is the maximum stack depth (i.e., the largest excess of `push` commands over `pop` commands) needed to process this file. Then comes  $t$ , the total number of pages (`bop` commands) present.

The postamble continues with font definitions, which are any number of `font.def` commands as described above, possibly interspersed with `nop` commands. Each font number that is used in the DVI file must be defined exactly twice: Once before it is first selected by a `font` command, and once in the postamble.

8. The last part of the postamble, following the *post.post* byte that signifies the end of the font definitions, contains *q*, a pointer to the *post* command that started the postamble. An identification byte, *i*, comes next; this currently equals 2, as in the preamble.

The *i* byte is followed by four or more bytes that are all equal to the decimal number 223 (i.e., '337 in octal). T<sub>E</sub>X puts out four to seven of these trailing bytes, until the total length of the file is a multiple of four bytes, since this works out best on machines that pack four bytes per word; but any number of 223's is allowed, as long as there are at least four of them. In effect, 223 is a sort of signature that is added at the very end.

This curious way to finish off a DVI file makes it feasible for DVI-reading programs to find the postamble first, on most computers, even though T<sub>E</sub>X wants to write the postamble last. Most operating systems permit random access to individual words or bytes of a file, so the DVI reader can start at the end and skip backwards over the 223's until finding the identification byte. Then it can back up four bytes, read *q*, and move to byte *q* of the file. This byte should, of course, contain the value 248 (*post*); now the postamble can be read, so the DVI reader discovers all the information needed for typesetting the pages. Note that it is also possible to skip through the DVI file at reasonably high speed to locate a particular page, if that proves desirable. This saves a lot of time, since DVI files used in production jobs tend to be large.

Unfortunately, however, standard PASCAL does not include the ability to access a random position in a file, or even to determine the length of a file. Almost all systems nowadays provide the necessary capabilities, so DVI format has been designed to work most efficiently with modern operating systems. But if DVI files have to be processed under the restrictions of standard PASCAL, one can simply read them from front to back, since the necessary header information is present in the preamble and in the font definitions. (The *l* and *u* and *s* and *t* parameters, which appear only in the postamble, are "frills" that are handy but not absolutely necessary.)