

NAME

pbm - portable bitmap file format

DESCRIPTION

The portable bitmap format is a lowest common denominator monochrome file format. It serves as the common language of a large family of bitmap conversion filters. Because the format pays no heed to efficiency, it is simple and general enough that one can easily develop programs to convert to and from just about any other graphics format, or to manipulate the image.

This is not a format that one would normally use to store a file or to transmit it to someone -- it's too expensive and not expressive enough for that. It's just an intermediary format. In its purest use, it lives only in a pipe between two other programs.

The format definition is as follows.

A PBM file consists of a sequence of one or more PBM images. There are no data, delimiters, or padding before, after, or between images.

Each PBM image consists of the following:

- A "magic number" for identifying the file type. A pbm image's magic number is the two characters "P4".
- Whitespace (blanks, TABs, CRs, LFs).
- The width in pixels of the image, formatted as ASCII characters in decimal.
- Whitespace.
- The height in pixels of the image, again in ASCII decimal.
- Newline or other single whitespace character.
- A raster of Height rows, in order from top to bottom. Each row is Width bits, packed 8 to a byte, with don't care bits to fill out the last byte in the row. Each bit represents a pixel: 1 is black, 0 is white. The order of the pixels is left to right. The order of their storage within each file byte is most significant bit to least significant bit. The order of the file bytes is from the beginning of the file toward the end of the file.
- Characters from a "#" to the next end-of-line, before the width/height line, are comments and are ignored.

There is actually another version of the PBM format, even more more simplistic, more lavishly wasteful of space than PBM, called Plain PBM. Plain PBM actually came first, but even its inventor couldn't stand its recklessly squanderous use of resources after a while and switched to what we now know as the regular PBM format. But Plain PBM is so redundant -- so overstated -- that it's virtually impossible to break. You can send it through the most liberal mail system (which was the original purpose of the PBM format) and it will arrive still readable. You can flip a dozen random bits and easily piece back together the original image. And we hardly need to define the format here, because you can decode it by inspection.

The difference is:

- There is exactly one image in a file.
- The "magic number" is "P1" instead of "P4".
- Each pixel in the raster is represented by a byte containing ASCII '1' or '0', representing black and white respectively. There are no fill bits at the end of a row.
- White space in the raster section is ignored.
- You can put any junk you want after the raster, if it starts with a white space character.
- No line should be longer than 70 characters.

Here is an example of a small bitmap in the plain PBM format:

```
P1
# feep.pbm
24 7
```

