

Tape formats for PDP-11 simulators

Jean-Marc Bourguet

17 october 2004

1 Introduction

When wanting to manipulate tapes for the PDP-10 on disk and on other systems, there are two problems to solve:

1. representing the tapes as file on disk. Tapes are record oriented devices (the transfer unit is a record of varying length) while files on current systems are not structured in record.
2. representing 36 bits data on 9 tracks tapes. The PDP-10 is a 36 bits computer and the tapes have 9 tracks used to store 8 bits bytes. Software on the PDP-10 used several methods to do the packing.

2 Representing tapes as file on disk

2.1 TPC

This is the simplest representation. The data for each record is preceeded by two bytes in the less significant byte first order representing the length of the record. If this length is odd, a padding byte (of undefined value) is added before the next record.

A tape mark (indication of end of file, two tape marks are used to mark the end of the recorded part of the tape) is an empty record.

2.2 TPE

In this representation, the record is preceeded by four bytes representing the length of the record (anew in little endian way). If the most significant bit of the header is 1, then there was an error in the record. The record is not padded but it is followed by a copy of the header so that simulating backward tape motion is easy.

2.3 TPS

This is nearly the same format as TPE, but there is a padding byte for odd length record.

3 Representing 36 bits data on 9 tracks tapes

3.1 Core-Dump

36 bits data is represented as 5 bytes. The first four contains the 32 most significant bits in most significant byte first order. The last byte contains the 4 less significant bits padded with 0 in the 4 most significant bits of the result.

B0	1	2	3	4	5	6	7	
	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23	
24	25	26	27	28	29	30	31	
0	0	0	0	32	33	34	35	

3.2 Industry-Compatible

The 32 most significant bits are represented as 4 bytes. The four remaining bits are ignored.

B0	1	2	3	4	5	6	7	
	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23	
24	25	26	27	28	29	30	31	

3.3 Sixbit

Groups of 6 bits are output.

0	0	B0	1	2	3	4	5
0	0	6	7	8	9	10	11
0	0	12	13	14	15	16	17
0	0	18	19	20	21	22	23
0	0	24	25	26	27	28	29
0	0	30	31	32	33	34	35

3.4 ANSI-ASCII

0	B0	1	2	3	4	5	6
0	7	8	9	10	11	12	13
0	14	15	16	17	18	19	20
0	21	22	23	24	25	26	27
35	28	29	30	31	32	33	34

With the less significant bit being 0 in ASCII files on the PDP-10, this allow easy transfer of text data.

3.5 High-density

B0	1	2	3	4	5	6	7
8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23
24	25	26	27	28	29	30	31
32	33	34	35	B0	1	2	3
4	5	6	7	8	9	10	11
12	13	14	15	16	17	18	19
20	21	22	23	24	25	26	27
28	29	30	31	32	33	34	35

This has an higher density and is also used by the FTP protocol (RFCxxx).

4 References

All the information in this document, and then some more is contained in the `vtape.txt` file in the `klh` distribution. I sadly found out this only after having examined the `simh` source and having experimented with setting SET TAPE FORMAT and examaning the result.