

File Name: CDBENCH.DOC

Date: 26.08.93

(c) T. Relph-Knight

Issue No. - 6

CDBENCH - C programs for testing CD-ROM drives

Introduction

The purpose of CDBENCH is to measure the performance of a CD-ROM drive and to save these measurements to a file. This benchmark is designed to treat a CD-ROM drive, its interface adapter (installed in a PC running MS-DOS) MSCDEX and the manufacturers device driver as an integrated sub-system. The benchmark tests the performance of this sub-system, the characteristics of the host PC should have no effect on the results.

CDBENCH is intended to give as clear a picture of CD-ROM drive performance as possible, to be comparable to the manufacturers claimed performance data and to be easily understood by the end user.

It does not try to simulate real world use as PCBENCH and the Windows Benchmarks do. CDBENCH is a specifically a sub-system test, simulation would only make sense if the CD-ROM drive were to be tested as an embedded part of a total system. In addition as CD-ROM is a read-only medium it is not possible to easily control and globally replicate the test data contents, neither are there a clearly identifiable group of commonly used applications to provide use profile data on which simulated application benchmarks could be based.

The software is not tied to any specific CD-ROM title, it will produce results with any ISO 9660 disc although for consistency it is recommended that the same disc title be used within any group of tests.

There are currently two programs in CDBENCH;

TRATECD.EXE and SEEKCD.EXE.

Trate measures the data transfer rate of the CD-ROM sub-system and Seekcd measures the seek times. Seekcd requires that the CD title used for testing contains at least 547 Mbytes of data to ensure the maximum seek takes place across the entire width of the disc. CDBENCH does not test for random access as CD-ROM is not by nature a random access medium.

Both programs output CSV format files of the tabulated results which can then be imported into any popular spreadsheet program, such as Excel, for formatting and printing in graphical form. Transfer rate and seek time plots from two different drives are included with this package as examples of the typical output of CDBENCH.

TRATECD

Tratecd measures the transmission rate of a CD-ROM sub-system by reading successive blocks of sectors which are incremented by one sector for each block read. The program starts by reading a block of one sector and increases this until a block of 200 sectors is read. Each block read is repeated a number of times (usually 10) and then averaged to produce a stable measurement. In tests it was found that, with standard rate drives (150 Kb/Sec.) an averaging value of ten was a good compromise between producing a clean graph curve and test duration. This should be doubled for each doubling in claimed transfer rate. A lower averaging constant produces results more quickly but the resultant graph has less clarity and accuracy. The C function 'clock()' is used to perform the timing and the benchmarks therefore rely on the accuracy of this function.

A maximum block size of 200 is used to allow the transfer rate curve to flatten out when the benchmark is used with the new multi-spin, double speed drives, which have a claimed maximum continuous transfer rate of 300 Kb/sec.. Most CD's have a sector size of 2Kbytes, which means that a block size of 150 sectors must be read, before the effect of read latency time is swamped. With the older single speed drives perhaps only 75 sectors need be read, as these drives claimed performance is only 150Kb/sec..

If the CD-ROM drive is unable to read the disc for any reason the program will exit with the message 'Not reading, is the disc loaded ?'.

A note on CD-ROM data rates

All CDs use a sector size of 2352 bytes. At standard audio spin rates, sectors are read at the rate of 75 sectors per second, resulting in a raw transfer rate of ; $2352 \times 75 = 176400$ bytes per second. For some reason some CD drive manufacturers quote this as their maximum transfer rate, converted to Kb/Sec. by dividing by 1000. Of the 2352 bytes 12 bytes are sync. bytes, 4 bytes are header bytes and 288 bytes are error correction bytes leaving 2048 bytes, or 2Kb as actual data. 2048 multiplied by 75 gives 153600 bytes per second, divided by 1000 this is 153.6 pseudo Kb/Sec., or by 1024 a real 150 Kbytes per second. Double and quad speed drives simply double or quadruple the sector read rate to 150 or 300 sectors per second.

The above explains why there are a number of different values quoted for the claimed drive transfer rates. In practice CD-ROM drives rarely reach the quoted sustained data rate due mainly to latency effects.

Latency

There is always a small delay between the point at which a CD-ROM drive locates the starting sector of a requested block of data and the point at which data transfer starts. Without latency, CD-ROM data transfer rate would be almost constant, for any block size from one sector upwards. With the small and constant latency delay, the data rate rises gradually, tending towards a constant rate but never quite reaching it. For small block sizes, 1 or 2 sectors, the latency delay is large, relative to the time taken to transfer data and therefore the calculated transfer rate is low. As block size is

increased this constant latency becomes less and less significant.

In fact latency is only relatively constant and does increase as the read head moves outwards towards the edge of the disk. Typical latency times are 50mS near the disk centre to 150mS near the edge, for a standard speed drive.

Latency can be measured by plotting transfer time against block size. This results in a stepped line, the steps are due to the effect of data buffering on the drive. The point where a line drawn through the bottom points of the steps crosses the vertical time axis is the latency time.

SEEKCD

Seekcd measures the positive and negative seek times of a CD-ROM drive at 16 equidistant points across a CD's written volume. This version of the code assumes a fixed volume size of 280000 sectors or 546.875 Mbytes. A volume size close to the maximum is used so that the read head moves over the largest possible distance for the maximum seek. This test will fail if a CD-ROM title with too small a data volume is used. If the CD-ROM drive is unable to read the disc for any reason the program will exit with the message 'Not reading, is the disc loaded ?'. For positive seeks the software returns the drive head to sector zero, before every seek, then measures the seek time for the head to reach the target sector, for negative seeks the head is moved to the reference sector, then the seek time to reach sector zero is measured. Each seek measurement is repeated and averaged to obtain a stable measurement.

If the CD-ROM used for the test does not have a volume of at least 280,000 sectors SEEKCD will exit with the error message;

'Maximum sector read failed, this CD-ROM is too small
Test needs a volume size of at least 280,000 sectors'

Running the Benchmarks

Having installed and tested a CD-ROM drive system and selected a disc with a large enough volume size.

Seekcd can be run by typing;

```
C:\seekcd <file>sk.csv
```

and Tratecd by typing;

```
C:\tratecd <file>tr.csv
```

at the prompt, as shown were <file> identifies the drive tested and the last two letters 'sk' or 'tr' identify the test run.

The results filename argument is optional but the results are saved only when the argument is present. The .CSV extension must be used for Excel to recognise the results file correctly. The result file is stored in CSV format, suitable for direct import into Excel (choose open from the file menu then set the file type to .CSV) or any other spreadsheet program, such as Supercalc. Once imported it then requires only a few minutes work to convert the tabulated results into a printed graph. (Two Excel macros SKCD.XLM and TRCD.XLM have been written which simplify and normalise this process).

Once the measurements have been completed and assuming a results file name was supplied at the command line, the program will request entry of the testers name and the name and/or model of CD-ROM drive being tested. These two input fields will accept an input string up to 60 characters in length.

For a direct print of the numeric results the program output may be redirected to a printer as follows;

```
C:\seekcd > lpt1   or   C:\tratecd > lpt1
```

Metrics

In order to produce an overall 'goodness' figure the tabular results from Tratecd and Seekcd will need to be interpreted and weighted. The exact method for this has still to be worked out but averages of the burst and sustained transfer rates and the average seek time will be weighted to provide an overall figure or figures.

Notes added 21.01.93

After some thought I propose that the metric is composed of four arithmetic averages. These are;

- 1/. Average of positive seek times.
- 2/. Average of negative seek times.
- 3/. Average transfer rate below a sector block size of 75 sectors.
- 4/. Average transfer rate above a sector block size of 75 sectors.

and the following timing,

- 5/. Trate test duration.

Further notes - 28.04.93

I have created a small Excel file as a CD-ROM results database and, using the first four values suggested above, first normalised each measurement, by taking the results for the Matsushita Kotobuki drive (a popular drive bundled with the Soundblaster card) and dividing each result by the Matsushita results (using reciprocals for the seek times). An overall performance is then calculated by averaging the normalised values. It may be that a more meaningful figure can be produced by introducing weighting values but this is not implemented at present.

Future Development

The current programs may be enhanced by adding further command line options for parameters such as the number of seek steps or the maximum numbers of sectors in a block. Version 1.0 has a dummy function to determine CD volume size as I have not been able to find a simple way of determining the volume. Hopefully this will be corrected in the next version.

Benchmark Proving Tests

To prove that the software behaves as expected and to provide guidelines for interpreting the results, Version 1.0 should be subjected to the following tests.

Results from at least two different CD-ROM titles run on the same drive system should be compared.

Results from at least two different drive systems, using the same host PC and CD title, should be compared.

Results from at least two different host PC's using the same drive and CD title should be compared.

Tests should be run to see whether the number of sector buffers assigned to MSCDEX have any affect on results.

Setting up the Host PC

The host PC should be set up with all unnecessary drivers, TSR etc. removed from the CONFIG.SYS and AUTOEXEC.BAT. The CONFIG.SYS must, of course, contain the command to load the manufacturers device driver for the CD-ROM drive and AUTOEXEC.BAT the command to load MSCDEX with suitable parameters.

Notes on test behavior

The original programs have now been run on a number of different drives including the NEC CDR-74 and Toshiba XM-3401 double speed drives and the Pioneer DRM-604X Mini-Changer quadruple speed drive.

These drives use caching and buffering techniques of greater sophistication than the previous second generation drives. It is obvious from these tests that the variable controlling the averaging loop, which repeats each measurement, should be a command line argument. For example with a loop variable of greater than one, the seek times for the Pioneer DRM-604X appeared to be 30 to 40 mSec. This was due to a sector caching on the first read, subsequent reads resulting in an apparent seek time of zero. The first measurement would be correct at around 340 mSec., averaging this out with nine zeroes gave apparent seek times of 34 mSec. Negative seek measurements are even more of a problem as at present they always seek to the same sector, sector zero and therefore fall foul of caching. Fixing this would require offsetting each negative seek measurement by a few sectors. For transmission rate measurements the average value needs to be increased, from 10 for the slower drives to perhaps 40 for the quad speed drives, to maintain as smooth a plot as possible.

It may also be useful to measure the total overall time taken to run the transfer rate test, including all the seek times and transfer times. At present the TRATECD program only measures the transfer times. A total time would give a good indication of a drives overall performance. With drives that have slow seek times, the effect of seek time on the overall test duration has been very noticeable, some drives take the best part of a day to run the whole test.

I have also noticed that many CD-ROM install programs allocate a value of 10 for the MSCDEX buffer option /M:. With the CDBENCH tests as they stand a value of 10 for MSCDEX buffers will limit the transfer rate performance. This appears as a flattening or distortion of the upper portion of the transfer rate curve. Increasing the buffer value to 20 or even 15 should correct this problem. For the purpose of a review it is probably best to check that /M: is always set to 20.

In some instances a CONFIG.SYS setting of FILES=20 has meant that results from SEEKCD have not been saved correctly. Increasing FILES to 30 should fix this problem.