

CDBench in action

PC Magazines benchmarks are always designed to present results as clearly as possible and so that results can be compared across the range of product reviewed. This inevitably means that although our benchmarks may produce thousands of individual measurements, this underlying detail is condensed into a single overall performance figure for each product.

Although we do not print all of our results, the fine detail these reveal is often telling and all our results are supplied to our reviewers. Our methods for reducing CDBench measurements to a single overall performance rating for each drive are described in the Labs Report section of this article.

CDBench results can be used to plot graphs of seek time and transfer rate and the shape of these graphs often reveals interesting qualities of particular drives.

For a hypothetical drive we will assume that a plot of the time taken to transfer blocks of data, which steadily increased in size, is a line with a constant slope, starting from an offset representing latency, as follows;

μ §

Time taken to read a block of data = (No. of sectors in a block x Time to read a block) + Latency

Re-plotting this, as transfer rate, results in the curve below, where for small transfers the fixed latency time is relative large, compared to the time it takes to actually transfer the data and so the measured transfer rate is low. As the volume of data increases the latency time has less and less effect and the transfer rate moves towards maximum.

μ §

In practice many real drives display the following characteristics;

μ s

This graph shows that the transfer performance for small blocks of data is largely affected by the drives read time latency. As the data block size read increases, read latency has less effect and the transfer rate gets closer and closer to maximum. A characteristic sawtooth ripple appears on this curve due to the data flowing from the buffer, until the buffer is almost empty, then data throughput falls while the buffer fills.

These two graphs are for the single speed type of drive, but the principle for the double and quad speed drives is the same.

Transfer efficiency is affected by the number of data buffers allocated by MSCDEX from system RAM. Unfortunately the number of buffers required can vary, depending on the nature of the application i.e. type of database, data type, CD-ROM directory structure, retrieval engine used etc.. Many CD-ROM drive installation routines set the buffers to 20Kb by setting the MSCDEX /M: command to 10. As sector data size is 2Kb, $10 \times 2 = 20\text{Kb}$. We found that, while running our benchmarks, the default buffer value was often insufficient to obtain the maximum transfer rates from the drives we tested. The following graph shows two transfer rate plots, one with the default value of 10 and one with the buffer parameter increased to 16.

μ §Clearly the transfer rate under our test conditions is much higher with /M: = 16. Increasing the buffer value beyond the a critical level will bring no further improvement and may even reduce overall application performance, as it will be tying up valuable system RAM.

Re-plotting a typical transfer rate curve as time against total data transferred shows the effect of data buffering more clearly, as in the following plot;

μ §

The stepped character of the above plot is typical of a buffered system, with the horizontal portion of the steps showing that the time taken to read data blocks that vary in size, within a certain range, is constant, while the rising part of the steps shows read time rising as the buffer is exhausted.

At the start of the curve we see a flattened section, where small blocks of data are transferred quickly, due to the burst capability of the SCSI bus.

Specific drives display interesting variations on the basic transfer rate curve, for example the following plot is for a double speed drive with a large buffer and a SCSI interface, which allows for rapid burst transfers, providing the block to be transferred is small.

μ §

Pioneers DRM604X quad speed drive produces a plot which shows evidence of intelligent caching to achieve the very high rates of which this drive is capable.

μ §

Further evidence emerged when we attempted to measure the Pioneer DRM-604X seek times. With our standard seek test we repeat each measurement ten times and average the result, to obtain a stable measurement. With the Pioneer drive we found this resulted in impossibly fast seek times, as subsequent seeks, following the first seek to any particular sector, were cached. Nine out of our ten measurements took only microseconds and reduced the overall average. Negative seek proved even more difficult to measure as, unlike positive seek where we were seeking to a fresh sector for each group of measurements, negative seeks always seek to sector zero. This meant our negative seek measurements were almost always close to zero, except when the cache flushed. Our next graph shows the seek results for the DRM-604X after we modified our software to make only one measurement to each seek point.

A more typical seek time plot, for a slower drive, is shown below;

μ s

In this plot the positive and negative seek times are quite separate, in others they may appear more closely intertwined.