

COMPARATIVE ANALYSIS OF HDD PERFORMANCE SOFTWARE AND THE EFFECTS OF USING A VIRTUAL HARD DISK

OBJECTIVE:

This lab is divided into two sections. Part one of this lab will set out to perform a comparative analysis of a variety of HDD performance software packages in order to determine some of the differences between each. The second part of this lab will perform the exact same experiments on a virtual hard disk (running from the same physical hard disk) in order to gain a better understanding of how a virtual disk performs as opposed to a hard disk. The lab will specifically seek to answer the following questions:

- How do the results vary between each software package?
- What might be the possible cause(s) for the differences?
- What affect does using a Virtual Disk have on each software package?
- Are the results from using a Virtual Disk *relatively* the same between each package? (Although the results are different, are the proportionally different between each package?)
- What might be the cause(s) for these differences?

EQUIPMENT/SOFTWARE USED:

Hardware:

- HP Pavilion f1703
- Western Digital Wd800B008 External Hard Drive

Software:

- Microsoft Windows XP Professional (x2)
- Microsoft Windows Virtual PC
- HD Tach
(<http://www.simplissoftware.com/Public/index.php?request=HdTach>)
- HD Speed (<http://www.snapfiles.com/get/hdspeed.html>)
- RoadKil DiskSpeed (<http://www.dirfile.com/freeware/drive.htm>)
- DiskSpeed32 (http://www.benchmarkhq.ru/english.html?/be_hdd.html)
- Drive Speed Test
(http://www.benchmarkhq.ru/english.html?/be_hdd.html)

PROCEDURES:

PART 1

The first portion of this lab involves the testing of a variety of HDD performance software packages. The experiments involved quantitative research. Each software package was compared against other software packages and evaluated according to the accuracy of the test (relative to performance median of all tests). The accuracy of each software package was determined by finding the median value of that specific test among all available software packages and comparing the specific package's results to that of the median value. Each software package was evaluated separately, and then a table containing all pertinent results was created to summarize the results of all tests.

For all tests, an external Western Digital Wd800B008 Hard Drive was used. This was done to eliminate some of the problems that may arise from conflicts with the Operating System attempting to retrieve information from the disk other than that being performed by the test. By testing a Hard Drive that is not associated with the Operating System, many of the risks of bias or contamination were removed.

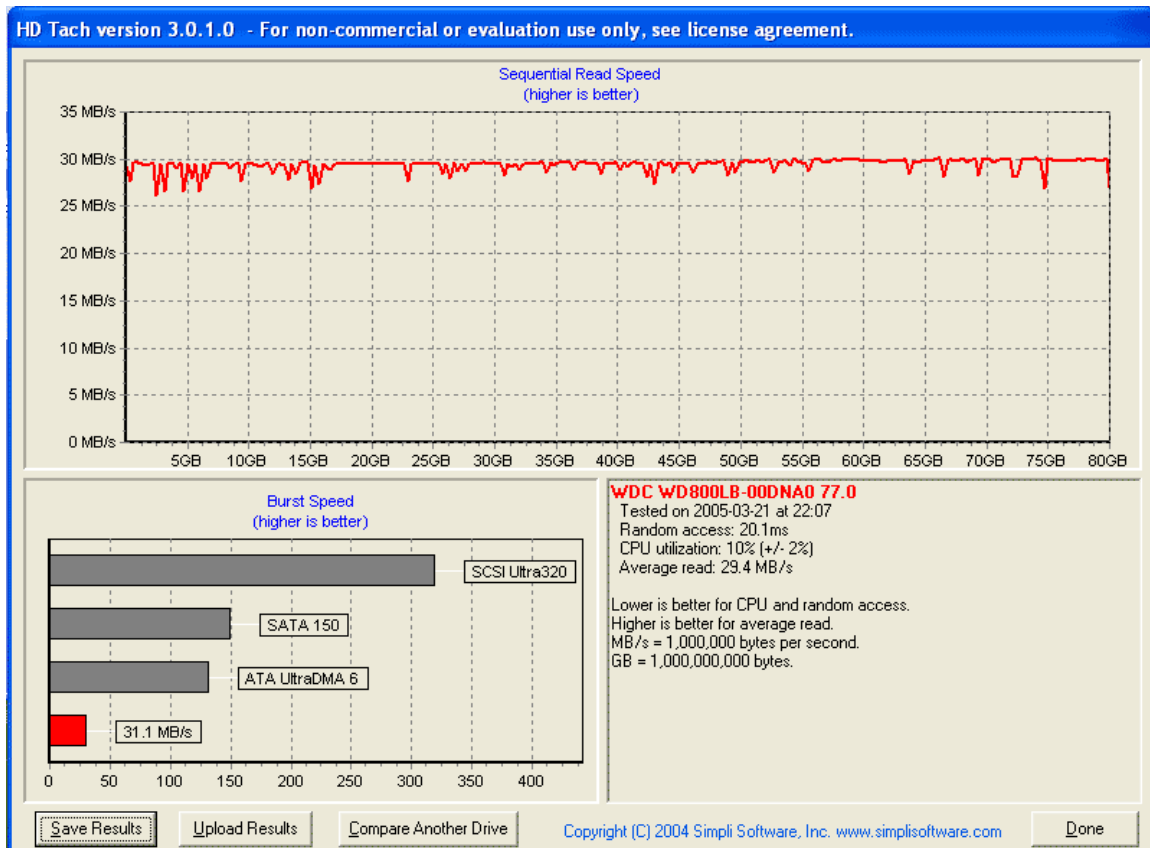
Some of the software packages could perform multiple tests. For the purposes of this lab, only sustained rate was taken into consideration, due to the fact that it was the only test offered by all software packages used in the comparison.

HD Tach

The first software package evaluated was *HD Tach*. This software provided both a "Quick bench (8MB)" and a "Long bench (32MB)". The software did not provide any information about what the tests actually did. The help documentation was also not useful, seeing as how the only information in the online help was the text "coming

soon...” The software also did not allow for exportation of the results into graphical format, but did provide a non-printable graph once the tests were complete. The resulting screenshot is shown below.

Figure 1 – HD Tach Results




From the graph, we can see that the both the burst and sustain rates are around 30 MB/s. The burst rate was 31.1 MB/s and the sustained rate was 29.4 MB/s. There were no particular anomalies associated with the performance of this experiment.

HD Speed

The next software tested was *HD Speed*. This software ran continually, and graphed the results in 'real-time' as the test progressed. Although there was an option to select "test burst rate", no change in output was apparent with this option selected.

Therefore, only the sustained rate was tested. This software also did not appear to have the option to export the results, thus a screenshot is provided below.

Figure 2 – HD Speed Results



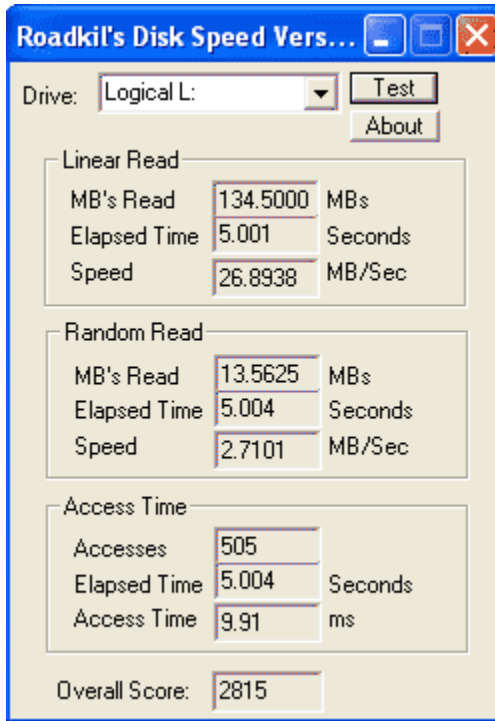
From the results of this test, we see that the average sustained read rate was 27.9 MB/s. There were two occasions during the test where the read time dropped dramatically. This may have been due to a problem with the software, a bad sector on the hard drive, or a higher priority process on the host Operating System.

RoadKil Disk Speed

The next software package that was tested was *RoadKil Disk Speed*. This product automatically performed three tests: Linear read (sustained), Random read, and Access time. The test then calculated an overall “score” based on all three tests.

However, no unit of measurement or comparison was provided, thus making the score virtually useless. A screenshot with the results are shown below.

Figure 3 – RoadKil Disk Speed Results



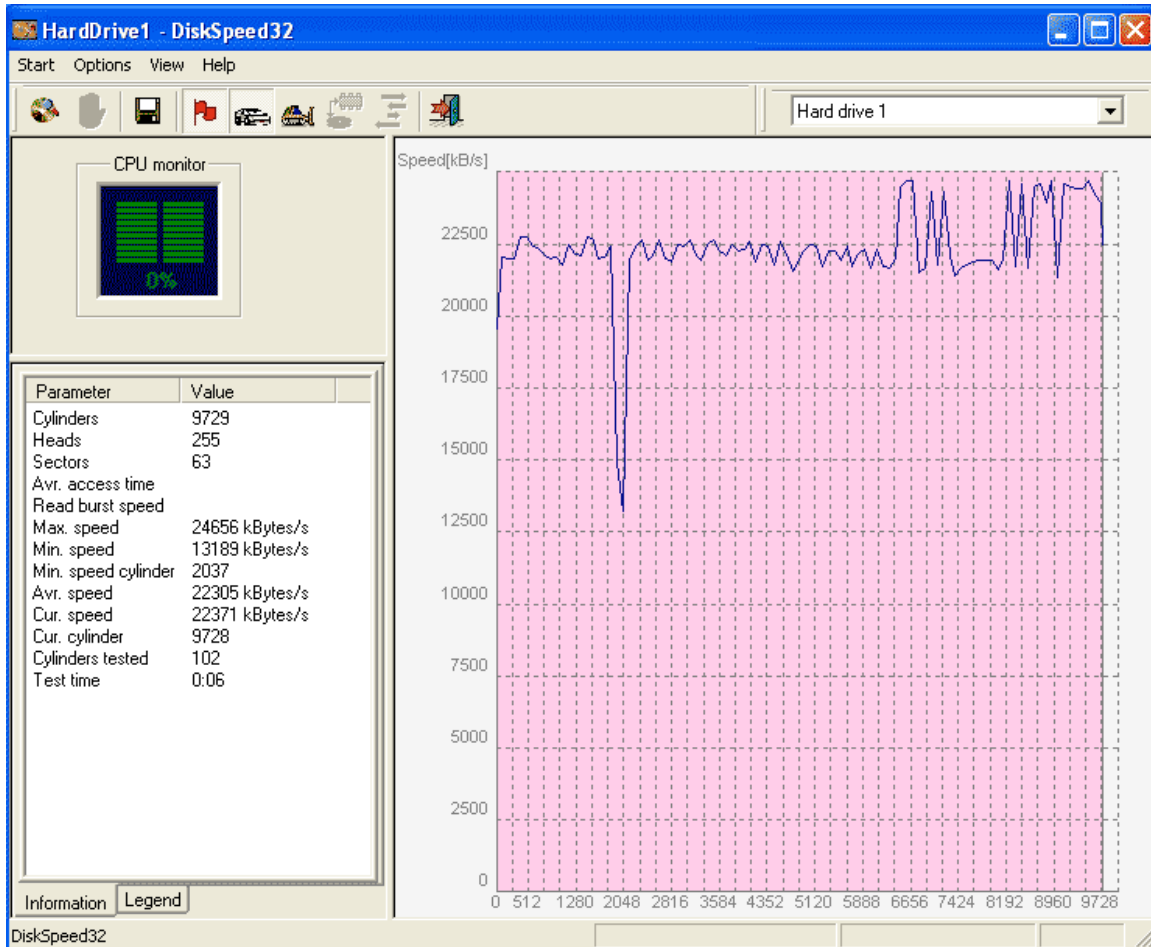
For the purposes of this lab, only the sequential read time is being evaluated. This software showed that it took approximately 5.001 seconds to read 134.50000 MB of data, resulting in an average read time of 26.8938 MB/s. The final number calculated by the software differed from an actual calculation which results in 26.8946 MB/s. However, this difference is negligible in light of the fact that the degree of accuracy of other software packages is only to one decimal place (in MB/s).

Disk Speed 32

The next software tested was *Disk Speed 32*. This software offered both an extended and a quick test. As the duration of most other software tests were relatively

short (several seconds to several minutes), the quick test for this application was selected. The results of this test are shown below.

Figure 4 – Disk Speed 32 Results

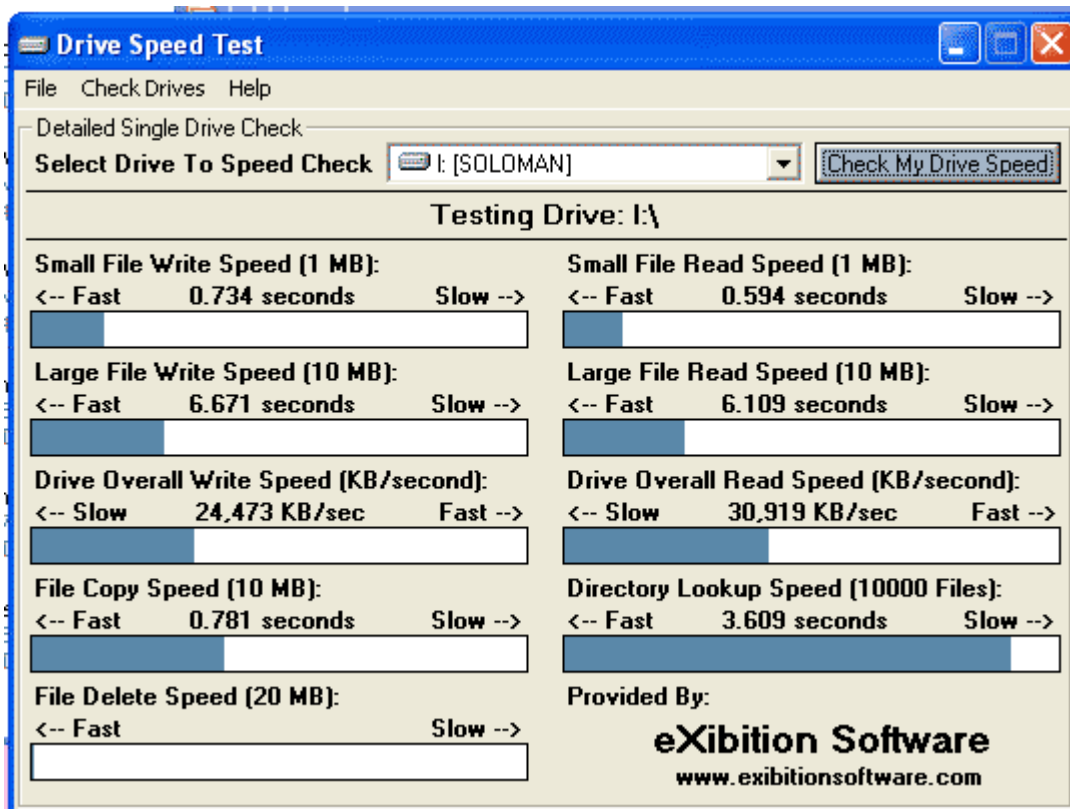


This test did not give a value for the burst speed. The maximum sustained speed was 24.7 MB/s while the average speed was 22.3 MB/s. There appeared to be one interruption or slow down during the process of the test. Also, towards the end of the test, the speed seemed to be increasing on average.

Drive Speed

The final software tested was *Drive Speed*. This software performed multiple tests incrementally, displaying the results for each test as it was completed. The final screenshot is shown below.

Figure 5 – Drive Speed Results



The test that appeared most similar to the other tests by the previous software packages was the "Drive Overall Read Speed." This resulted in an average speed of 30.92 MB/s. Interestingly, this was the only software package among those tested that also tested the performance of writing to the hard drive. The other software packages performed a variety of tests, but were all relating to reading data as opposed to writing data.

The summary of results and analysis of the results can be found in the “Report” section of this document.

PART 2

The second portion of this lab consisted of running the same tests as previously, but this time running the tests on a Virtual Hard Disk. In order to better understand how a Virtual Hard Disk differs from a physical hard disk, Virtual PC (for Windows) was installed on the same machine as was used for the previous tests. The same version of MS Windows XP Professional (with a different license) was installed as the subsequent virtual operating system. The virtual operating system was saved on the same external disk drive as was used for the previous tests. This ensured that the same physical drive was being tested by both the physical tests and the virtual tests.

Because the testing process is identical as the previous portion of this lab, only the resulting screenshots of each experiment will be inserted into this report, with a summary of observations at the end of this section of the document. A complete analysis of the results will be discussed in the “Report” section of this document.

Figure 6 – HD Tach Restuts (Virtual)




Figure 7 – HD Speed Results (Virtual)




Figure 8 – Roadkil Disk Speed Results (Virtual)

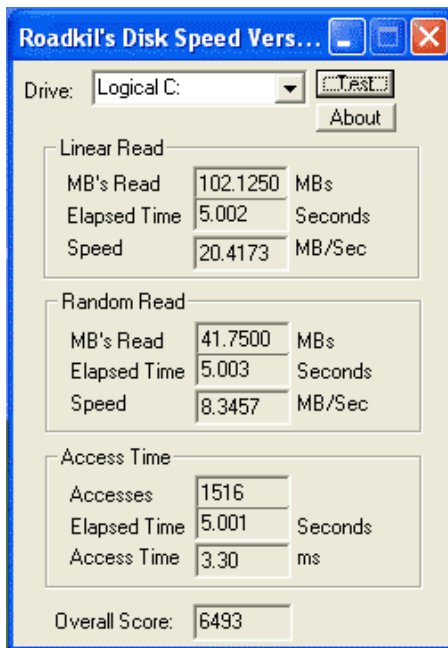
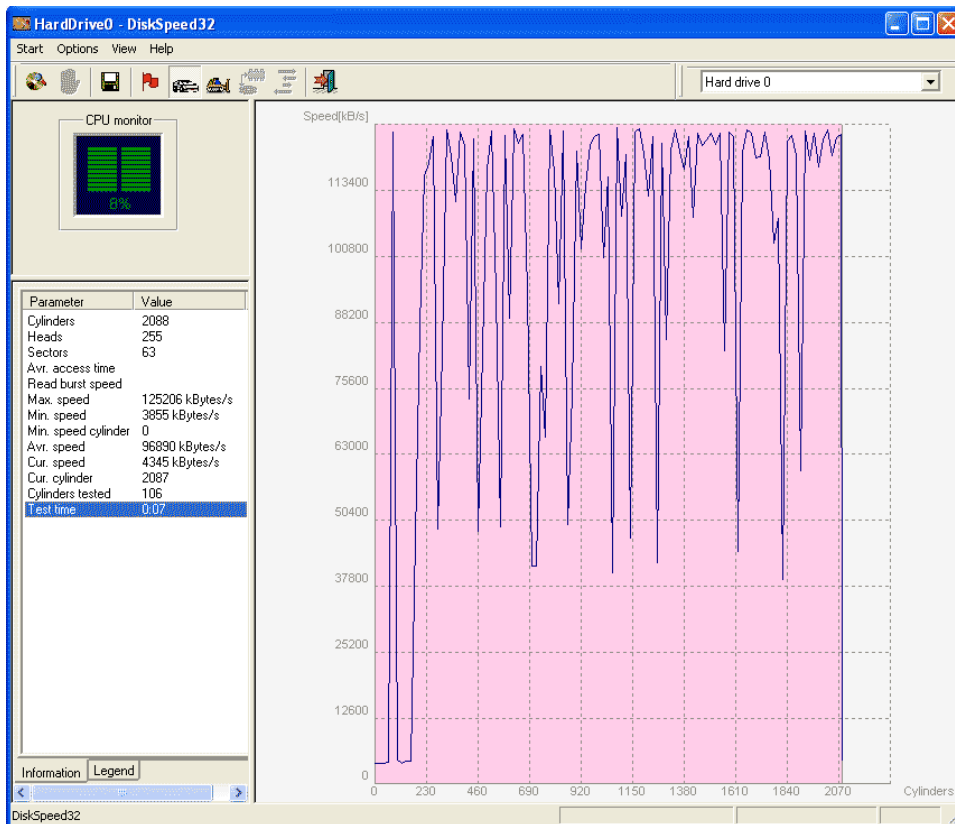


Figure 9 – Disk Speed 32 Results (Virtual)



It should first be noted that there are no results to display for the *Drive Speed* software because an error was produced when attempting to run the software. The error was undefined and although several attempts were made to get the software to run, it was never successfully able to perform a test on the Virtual Hard Disk.

As a general observation, in two of the three tests that produce graphical results, the read time seems quite erratic when running on a virtual hard disk. Also, the speed, overall, is much greater when running on a virtual disk than when running from the hard disk. Whether or not this would stand true through an extended test of the entire hard drive would be a suggestion for future research.

The RoadKil software actually indicated a slower read time, which was quite different than the faster times indicated by all of the other software packages.

REPORT:

Both portions of this test yielded interesting results. The subsequent result tables are summarized below. Related graphs are also provided.

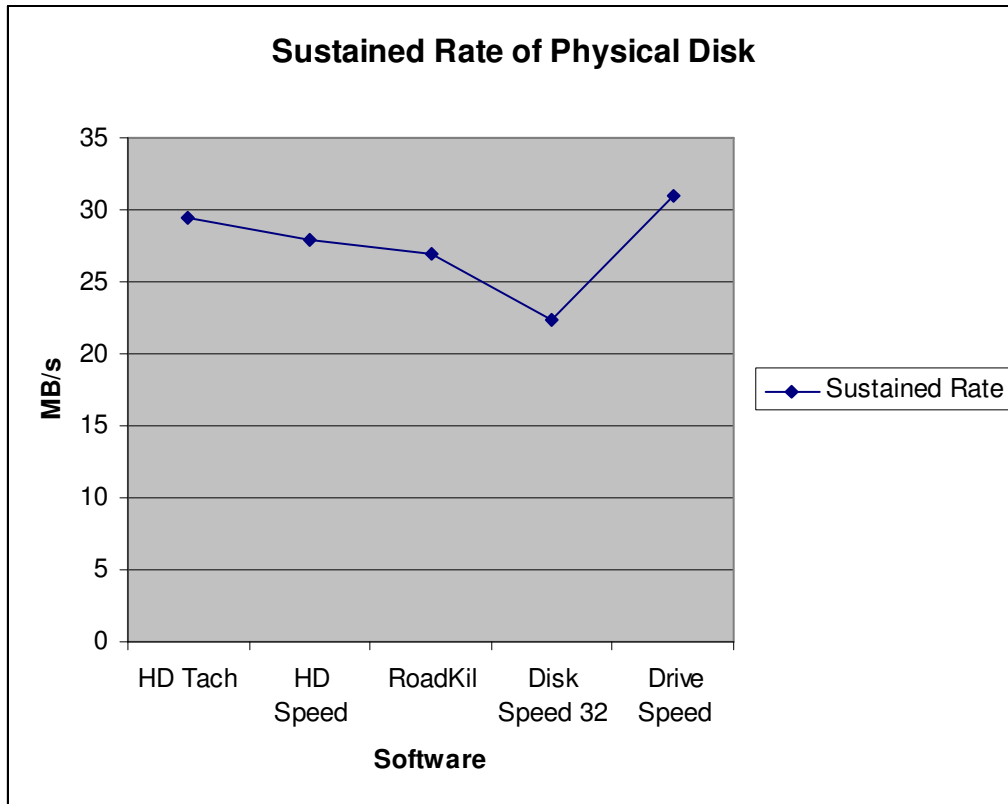
Table 1 – Summary for Part 1 (Physical Disk)

PART 1					
Software	HD Tach	HD Speed	RoadKil	Disk Speed 32	Drive Speed
Sustained Rate	29.4	27.9	26.9	22.3	30.92

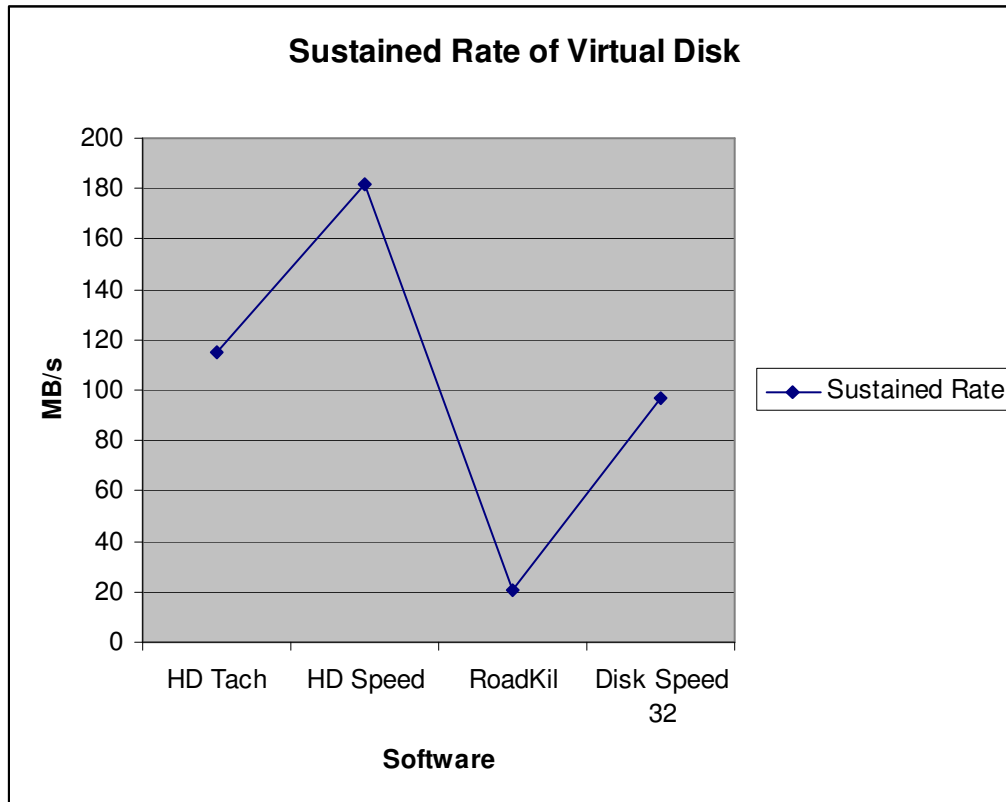
Table 2 – Summary for Part 2 (Virtual Disk)

PART 2					
Software	HD Tach	HD Speed	RoadKil	Disk Speed 32	Drive Speed
Sustained Rate	115	182.1	20.4	96.9	NA

Graph 1 – Sustained Rates for Physical Disk



Graph 2 – Sustained Rates for Virtual Disk



The median value for the tests performed on the physical disk is 27.9 MB/s. The median value for the tests performed on the virtual disk is 106 MB/s. From this, we can see that the overall speed of the virtual disk was greater than that of the physical disk. However, general observations of the data and charts indicate that the tests performed on the Virtual Disk were erratic and inconsistent. The rate of the physical disk was fairly consistent.

When comparing the software packages from part 1, HD Speed was the median of all of the software packages, thus deviation from that point can be used to make a relative comparison of the software packages. Disk Speed 32 provided a slower reading rate than the others, while Drive Speed provided a higher rate than the others. When

comparing the results from part 2, it is difficult to make any assumptions about the data due to the wide variance and erratic circumstances that the data present.

Upon completing the experiments, an important observation was made concerning the sustained rate of the physical disk. The physical disk was an external hard drive connected via a USB 2.0 port. For this reason, the rate seems to be limited by the transfer rate of the connecting device. This provides a good example of how the interface becomes important when considering storage devices. This also helped to indicate that although the virtual operating system (and thus virtual hard drive) were running from the same external hard drive, in order to achieve the rates that the tests indicated, some sort of local disk virtualization must be taking place within the virtual software. For example, some sort of transfer buffer, or temporary memory source may be used while reading from, or writing to the virtual disk. This may also help to explain how erratic the results appeared to be. For example, if a temporary buffer is used, then the local memory would appear faster, but while then sending that buffered data to the actual disk, the rate would slow down and again be limited by the rate of the USB 2.0 interface.

CONCLUSIONS:

The results of this lab can be considered a success in light of the fact that each objective of the lab was met. Consider the following questions:

- How do the results vary between each software package?
 - For part 1, the variance was relatively small, ranging from 22.3 MB/s to 30.9 MB/s. The median for part one was 27.9 MB/s. For part two, the range of variance was quite extreme, ranging from 20.4 MB/s to 182.1 MB/s. Although the median was calculated to be 106 MB/s, this information is deemed unreliable due to the erratic nature of the tests and resulting data.
- What might be the possible cause(s) for the differences?
 - Part 1 appears to be limited by the interface (USB 2.0) rate, and thus a suggestion for future study would be to perform the same tests on a secondary hard drive that is connected through IDE or a similar interface. In spite of this, the variance between the different software packages is not easily explained except by speculation. The software packages did not provide technical information as to how the tests were performed, however, it can be speculated that the differences had something to do with the way that the tests were performed internally. It is possible that one software package may use inefficient coding practices that slow down the processing, thus resulting in slightly skewed data. It could also have to do with what portion of the HDD was being analyzed at the time of the test. If there was a bad sector, then that may have delayed the tests slightly. It may also be a result of miscellaneous activity taking place on the operating system.
- What affect does using a Virtual Disk have on each software package?
 - Using the virtual disk resulted in erratic and erroneous results. The actual graphs produced by the performance software indicated an inconsistency in the performance of the virtual disk. This may be explained by virtualization taking place between the host operating system and the virtual system. If buffering of data transfer is used, then it might result in data similar to the data produced from these tests.
- Are the results from using a Virtual Disk *relatively* the same between each package? (Although the results are different, are the proportionally different between each package?)
 - No. The results obtained from using the virtual disk were not relatively or proportionally different than those from the physical disk. This indicates that the virtual disk does NOT transfer data directly to the physical disk.
- What might be the cause(s) for these differences?

- Once again, these differences are most likely caused by virtualization or buffering of the hard disk. The virtual operating system apparently does not directly transfer data to and from the underlying physical disk. It is possible that some sort of temporary storage is being used. This might be a buffer in RAM or a temporary buffer on the local disk of the host system.

Overall, this lab was a valuable learning experience. The importance of the storage interface speed became more apparent, as the speed of the physical disk seemed limited by the USB 2.0 interface. Also, a deeper knowledge of how virtual storage differs from physical storage was also an important result of this lab. Although the results of some of the experiments were not as expected, the objectives of the lab were met and thus the lab concluded successfully.