Automated Disk Drive Characterization

Jiri Schindler, Gregory R. Ganger
Carnegie Mellon University
Pittsburgh, PA 15213
{schindjr, ganger}@ece.cmu.edu

ABSTRACT

DIXtrac is a program that automatically characterizes the performance of modern disk drives. This extended abstract overviews the contents of [3], which describes and validates DIXtrac's algorithms for extracting accurate values for over 100 performance-critical parameters in 2 – 6 minutes without human intervention or special hardware support. The extracted data includes detailed layout and geometry information, mechanical timings, cache management policies, and command processing overheads. DIXtrac is validated by configuring a detailed disk simulator with its extracted parameters; in most cases, the resulting accuracies match those of the most accurate disk simulators reported in the literature. To date, DIXtrac has been successfully used on ten different models from four different manufacturers. A growing database of validated disk characteristics is available in DiskSim [1] format at `http://www.ece.cmu.edu/~ganger/disksim/diskspecs.html`.

1. INTRODUCTION

While detailed models of disk behavior are highly desirable for aggressive scheduling algorithms and comprehensive system modeling, obtaining accurate data from state-of-the-art disk drives for detailed disk models is at best tedious. Although methodologies for accurate disk models are well understood and general techniques for on-line extraction of parameters from SCSI disk drives have been proposed, there is no available system for easily retrieving parameters from a given disk drive.

This paper presents DIXtrac (DIsk eXtraction), a program that can quickly and automatically characterize disk drives that understand Small Computer System Interface (SCSI). Without human intervention, DIXtrac can discover accurate values for over 100 performance-critical parameters. It runs on standard Linux 2.2 systems and requires no special hardware or operating system support. By automating this process, DIXtrac greatly simplifies the process of collecting disk drive characterizations.

Automatic extraction of parameters from disks poses several challenges compared to manual, or user-assisted, approaches. If the extraction is to be general enough to work for disk drives of different vendors, it must use methods that will work for all drives and dynamically adapt when extracting from disks that support only a subset of the interface features. DIXtrac's extraction algorithms use only widely-supported interface features and include multiple approaches for discovering many characteristics.

2. CHARACTERIZING DISK DRIVES

To completely characterize a disk drive, one must describe the disk's geometry and layout, mechanical timings, cache parameters and behavior, and all command processing overheads. Thus, the characterization of a disk consists of a list of performance-critical parameters and their values. Naturally, such a characterization makes implicit assumptions about the general functionality of a disk.

Therefore, to accurately characterize a disk drive, there exists a set of requirements that the disk interface must meet. First, it must be possible to determine disk's geometry either experimentally or from manufacturer's data. Second, it must be possible to read and write specific physical locations. Also, while not strictly necessary, it is very useful to be able to temporarily turn off the on-disk cache. With just these capabilities, DIXtrac can determine over 100 performance-critical parameters expected by the proven DiskSim simulator [1].

The first requirement enables construction of a detailed disk map that identifies the physical location of each logical block number (LBN) exposed by the disk interface. Once a disk map is created, the second requirement enables the assembly of appropriate test vectors consisting of read and write requests. By timing the completion times of individual requests in the test vector, the various parameters can be obtained as first described in [4] and refined and automated by this work.

DIXtrac currently works on SCSI disks as follows. First, the Translate option of the SEND DIAGNOSTIC and RECEIVE DIAGNOSTIC commands translates a given LBN to its physical address on the disk. The READ DEFECT LIST command gives the physical locations of all defective sectors. With these two commands, DIXtrac can create a complete and accurate disk map. Second, the READ and WRITE commands take a starting LBN and a number of consecutive blocks to be read or written. Third, the cache can usually be enabled and disabled by changing the Cache Mode Page with the MODE SELECT command.

Vendor	IBM	Quantum
Disk Model	Ultrastar 18ES	Atlas 10K
Model Number	DNES309170	TN09100W
Task	*Time (seconds)*	
Layout extraction	164.7 (10.6)	50.1 (3.9)
Complete seek curve	45.2 (0.1)	33.3 (0.3)
Mechanical overheads	35.8 (1.3)	18.6 (1.5)
Cache parameters	25.6 (0.6)	12.6 (0.3)
Cmd process overheads	64.3 (2.5)	12.7 (0.9)
Totals	**335.6 (9.2)**	**127.4 (4)**
Capacity (sectors)	17916239	17783248
Defects	123	64
One LBN Translation	2.41 ms	0.86 ms
Translations	36911 (39)	26437 (79)

Table 1: Break down of extraction times and address translation characteristics for some tested disks. The times are mean values of five extractions. The values in parentheses are standard deviations.

3. RESULTS AND PERFORMANCE

DIXtrac has been fully tested on ten disk models: IBM Ultrastar 18ES, Hewlett-Packard C2247, Quantum Viking, Quantum Atlas III, Quantum Atlas 10K, Seagate Barracuda 4LP, Seagate Cheetah 4LP, Seagate Cheetah 9LP, Seagate Cheetah 18LP, and Seagate Hawk. For brevity, this section lists results for two 9.1 GB disks: IBM Ultrastar 18ES (7200 RPM 1998 disk) and Quantum Atlas 10K (10000 RPM 1999 disk).

3.1 Extraction Times

Table 1 summarizes the extraction times. The times are broken down to show how long each extraction step takes. For most disks an entire characterization takes less than three minutes. Further reduction of run time is possible, at the expense of accuracy, by using fewer repetitions for the timing extractions.

The table also shows the number of address translations required by DIXtrac to characterize each disk. Note that the number of translations does not depend directly on the capacity of the disk. Instead, it depends mainly on the sparing scheme, the number of zones, and the number of defects. Comparing the number of sectors to the number of translations provides a metric of efficiency for DIXtrac's layout discovery algorithms. For example, translation of every logical block ranges from almost 5 hours (Atlas 10K) to 21 hours (Barracuda 4LP).

3.2 Validation of Extracted Values

After extracting parameters from each disk drive, a synthetic trace was generated, and the response time of each request in the trace was measured on the real disk. The trace run and the extracted parameters were then fed to DiskSim to produce simulated per-request response times. The real and simulated response times are then compared.

Disk Model	Ultrastar 18ES		Atlas 10K	
Trace Type	mixed	random	mixed	random
$RMS_{Overall}$ (ms)	**0.20**	**0.07**	**0.30**	**0.19**
% $\overline{T}_{Mean}$	3.9%	0.6%	8.6%	2.2%

Table 2: Demetit figures. $RMS_{Overall}$ is the overall demerit figure for 25 trace runs combined. The % $\overline{T}_{Mean}$ value is the percent difference of the respective RMS from the mean real disk response time.

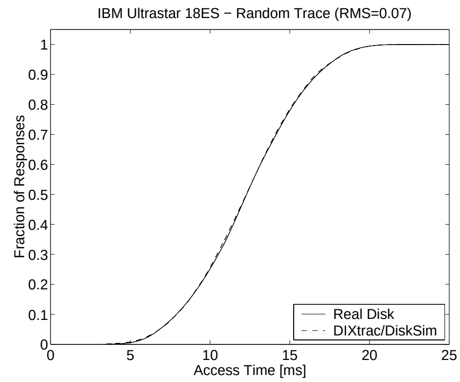

Figure 1: Measured and simulated response time cumulative distribution for IBM Ultrastar 18ES disk.

Two synthetic traces were used to test the extracted parameters. Both traces consist of 5000 requests with 2/3 READs and 1/3 WRITEs. The requests of the first trace are uniformly distributed across the entire disk drive. The second trace focuses on the cache behavior and was executed with a mix of 20% sequential requests and 30% local (within 500 LBNs) requests.

The difference between the response times of the real disk and the DIXtrac-configured simulator can be quantified by a demerit figure [2], which is the root mean square distance between the two curves. The demerit figure, here referred to as the RMS, for each disk is given in Table 2. The obtained demerit figures compare favorably with the most accurate disk simulation models reported in the literature [2; 4].

4. ACKNOWLEDGEMENTS

This work was supported in part by the NSF under grant #ECD-8907068. The authors are indebted to generous contributions from the member companies of the Parallel Data Consortium and Data Storage Systems Center at Carnegie Mellon Univeristy.

5. REFERENCES

[1] G. R. Ganger, B. L. Worthington, and Y. N. Patt. The DiskSim simulation environment. Technical Report CSE-TR-358-98, Dept. of Electrical Engineering and Computer Science, Univ. of Michigan, February 1998.

[2] C. Ruemmler and J. Wilkes. An introduction to disk drive modeling. *IEEE Computer*, 27(3):17–28, March 1994.

[3] J. Schindler and G. R. Ganger. Automated disk drive characterization. Technical Report CMU-CS-99-176, Carnegie Mellon University, December 1999.

[4] B. L. Worthington, G. R. Ganger, and Y. N. Patt. On-line extraction of SCSI disk drive parameters. In *Proc. of ACM Sigmetrics Conference*, pages 146–156, 1995.