A Nine Year Study of File System and Storage Benchmarking

AVISHAY TRAEGER and EREZ ZADOK
Stony Brook University
and
NIKOLAI JOUKOV and CHARLES P. WRIGHT
IBM T.J. Watson Research Center

Technical Report FSL-07-01

Benchmarking is critical when evaluating performance, but is especially difficult for file and storage systems. Complex interactions between I/O devices, caches, kernel daemons, and other OS components result in behavior that is rather difficult to analyze. Moreover, systems have different features and optimizations, so no single benchmark is always suitable. The large variety of workloads that these systems experience in the real world also add to this difficulty.

In this article we survey 415 file system and storage benchmarks from 106 recent papers. We found that most popular benchmarks are flawed and many research papers do not provide a clear indication of true performance. We provide guidelines that we hope will improve future performance evaluations. To show how some widely-used benchmarks can conceal or over-emphasize overheads, we conducted a set of experiments. As a specific example, slowing down read operations on ext2 by a factor of 32 resulted in only a 2–5% wall-clock slowdown in a popular compile benchmark. Finally, we discuss future work to improve file system and storage benchmarking.

Categories and Subject Descriptors: D.4.8 [**Operating Systems**]: Performance—*Measurements*; C.4 [**Performance of Systems**]: Measurement Techniques and Performance Attributes

General Terms: Design, Experimentation, Measurement, Performance

Additional Key Words and Phrases: Benchmarks, File Systems, Storage Systems

1. INTRODUCTION

Benchmarks are most often used to provide an idea of how fast some piece of software or hardware runs. The results can significantly add to, or detract from, the value of a product (be it monetary or otherwise). For example, they may be used by potential consumers in purchasing decisions or by researchers to help determine a system's worth.

When a performance evaluation of a system is presented, the results and implications must be clear to the reader. This includes accurate depictions of behavior under realistic workloads and in worst-case scenarios, as well as explaining the reasoning behind bench-

Author's Address: Avishay Traeger, Computer Science Dept, Stony Brook University, Stony Brook, NY 11794

marking methodologies. In addition, the reader should be able to verify the benchmark results, and compare the performance of one system with that of another. To accomplish these goals, much thought must go into choosing suitable benchmarks and configurations, and accurate results must be conveyed.

Ideally, users could test performance in their own settings using real workloads. This transfers the responsibility of benchmarking from the author to the user. However, this is usually impractical because testing multiple systems is time consuming, especially because exposing the system to real workloads implies learning how to configure the system properly, possibly migrating data and other settings to the new systems, as well as dealing with their respective bugs. In addition, many systems developed for research purposes are not released to the public. Although rare, we have seen performance measured using actual workloads when they are created for in-house use [Ghemawat et al. 2003] or are made by a company to be deployed [Schmuck and Haskin 2002; Eisler et al. 2007]. The next best thing is for some party (usually the authors) to run workloads that are representative of real-world use on commodity hardware. These workloads come in the form of synthetic benchmarks, executing real programs, or using traces of some activity. Simulating workloads raises concerns about how accurately these benchmarks portray the end-user's workload. Because of this, benchmarks must be well-understood so as to not have unknown side-effects, and should provide a good approximation of how the program would perform under different loads.

Benchmarking file and storage systems requires extra care, which exacerbates the situation. Even though these systems share the goal of providing access to data via a uniform API, they differ in many ways, such as the type of underlying media (e.g., magnetic disk, network storage, CD-ROM, volatile RAM, flash RAM, etc.), the storage environment (e.g., RAID, LVM, virtualization, etc.), the workloads that the system is optimized for, and their features (e.g., journals, encryption, etc.).

In addition, complex interactions exist between file systems, I/O devices, specialized caches (e.g, buffer cache, disk cache), kernel daemons (e.g., `kflushd`), and other OS components. Some operations may be performed asynchronously, and this activity is not always captured in benchmark results. Because of this complexity, many factors must be taken into account when performing benchmarks and analyzing results.

In this article we concentrate on file and storage system benchmarks in the research community. Specifically, we comment on how to choose and create benchmarks, how to run them, and how to analyze and report the results. We have surveyed a selection of recent file and storage system papers, and have found several poor benchmarking practices, as well as some good practices. We classify the benchmarks into three categories and discuss them in turn:

Macro-benchmarks. The performance is tested against a particular workload that is meant to represent some real-world workload.

Trace replays. A program replays operations which were recorded in a real scenario, with the hope that it is representative of real-world workloads.

Micro-benchmarks. Few (typically one or two) operations are tested to isolate their specific overheads within the system.

The rest of this article is organized as follows. In Section 2 we describe the criteria for selecting publications for our survey and list the papers we analyze. Section 3 provides suggested guidelines to use when benchmarking and Section 4 discusses how well the surveyed papers have followed those suggestions. In Section 5 we give an overview of related research. We describe the system configuration and benchmarking procedures that were used in our experiments in Section 6.

Section 7 reviews the pros and cons of the macro-benchmarks that were used in the surveyed papers; we also include a few other notable benchmarks for completeness. We identified problems that can cause results to be inaccurate, misleading, incomparable, or unreproducible. Additionally, the workloads used for benchmarking may have little to do with workloads that the system will need to handle in the real world. Readers should be aware of the shortcomings so they can know how to interpret the results correctly.

In Section 8 we examine how the papers that we surveyed used traces and describe four main problems that arise when using traces for performance analysis. First, the methods used to capture traces are not always specified and this can affect how the results should be interpreted. Second, the methods used to replay traces may not be accurate because they generate workloads that are different from the traced workload. Third, the traces may not be representative of the real-world environment that the researcher is aiming to capture if the trace is too short. Fourth, trace workloads are reproducible only as long as other researchers have access to the traces.

Section 9 describes the widely-available micro-benchmarks in the same way as we describe the macro-benchmarks. Since custom, or ad-hoc, micro-benchmarks are of little interest on their own, we show a selection of examples that illustrate good and bad ways to utilize these benchmarks. Micro-benchmarks are useful to isolate the performance of parts of the system because the benchmarks do not have the added complications that arise from exercising several operations at once. Although micro-benchmarks provide the most fine-grained information, they do not usually provide enough information about the overall performance of a system, and even the results from several different micro-benchmarks can leave a picture incomplete.

In Section 10 we discuss some of the more popular programs that can generate workloads according to some specifications that the user provides. Workload generators are generally less flexible than custom benchmarks, but there is no need to create a program from scratch, which would be a less reproducible solution.

Section 11 describes a suite of tools for benchmarking automation, which can save time and avoid errors associated with repetitive tasks. Section 12 shows the benchmarks that we performed. We describe the file system that we used to conduct the tests and the experiments themselves. The experiments show how benchmarks can hide overheads. We conclude in Section 13 and give a summary of suggestions for choosing the proper benchmark, and offer our ideas for the future of file and storage system benchmarking.

2. SURVEYED PAPERS

Research papers have used A variety of benchmarks to analyze the performance of file and storage systems. This paper surveys the benchmarks and benchmarking practices from a selection of the following recent conferences:

—the Symposium on Operating Systems Principles (SOSP 1999, 2001, 2003, and 2005)

—the Symposium on Operating Systems Design and Implementation (OSDI 2000, 2002, 2004, and 2006)

—the USENIX Conference on File and Storage Technologies (FAST 2002, 2003, 2004, 2005, and 2007)

—the USENIX Annual Technical Conference (USENIX 1999, 2000, 2001, 2002, 2003, 2004, 2005, and 2006)

Research papers relating to file systems and storage often appear in the proceedings of these conferences which are considered to be of high quality. We decided to consider only full-length papers from conferences that have run for at least five years. In addition, we consider only file systems and storage papers that have evaluated their implementations (no simulations), and from those, only benchmarks that are used to measure performance in terms of latency or throughput. For example, benchmarks used to verify correctness or report on the amount of disk space used were not included. Studies that are similar to ours were performed in the past [Small et al. 1997; Mogul 1999], and so we believe that this cross section of conference papers is adequate to make some generalizations.

We surveyed 106 papers in total, eight of which are our own [Lu et al. 2002; Zhang et al. 2002; Memik et al. 2002; Lumb et al. 2002; Schmuck and Haskin 2002; Ng et al. 2002; Anderson et al. 2002; Sobti et al. 2002; Kim et al. 2002; Patterson et al. 2002; Miller et al. 2002; Schindler et al. 2002; Quinlan and Dorward 2002; Aguilera et al. 2003; Flinn et al. 2003; Dimitrijevic et al. 2003; Lumb et al. 2003; Magoutis et al. 2003; Soules et al. 2003; Kallahalla et al. 2003; Rhea et al. 2003; Sivathanu et al. 2003; Sarkar et al. 2003; DeBergalis et al. 2003; Zhang and Ghose 2003; Thereska et al. 2004; Radkov et al. 2004; Muniswamy-Reddy et al. 2004; Schindler et al. 2004; Sivathanu et al. 2004; Tolia et al. 2004; Corbett et al. 2004; Aranya et al. 2004; Joglekar et al. 2005; Joukov et al. 2005; Sivathanu et al. 2005; Denehy et al. 2005; Schlosser et al. 2005; Vilayannur et al. 2005; Peterson et al. 2005; Abd-El-Malek et al. 2005; Wachs et al. 2007; Veeraraghavan et al. 2007; Eisler et al. 2007; Peterson et al. 2007; Gulati et al. 2007; Weddle et al. 2007; Tian et al. 2007; Wang and Merchant 2007; Yumerefendi and Chase 2007; Cipar et al. 2007; Gopal and Manber 1999; Grönvall et al. 1999; Wang et al. 1999; Kim et al. 2000; Fu et al. 2000; Anderson et al. 2000; Meter and Gao 2000; Strunk et al. 2000; Yu et al. 2000; Adya et al. 2002; Muthitacharoen et al. 2002; Saito et al. 2002; MacCormick et al. 2004; Nightingale and Flinn 2004; Sivathanu et al. 2004; Gniady et al. 2004; Li et al. 2004; Weil et al. 2006; Peek and Flinn 2006; Nightingale et al. 2006; Sivathanu et al. 2006; Santry et al. 1999; Mazieres et al. 1999; Muthitacharoen et al. 2001; Rowstron and Druschel 2001; Dabek et al. 2001; Kaminsky et al. 2003; Ghemawat et al. 2003; Arpaci-Dusseau et al. 2003; Huang et al. 2005; Zhu et al. 2005; Prabhakaran et al. 2005; Nightingale et al. 2005; Zadok et al. 1999; Lee et al. 1999; Zadok and Nieh 2000; Seltzer et al. 2000; Maziéres 2001; Huang and Chiueh 2001; Kroeger and Long 2001; Zadok et al. 2001; Zhou et al. 2001; Stein et al. 2001; Wang et al. 2002; Denehy et al. 2002; Magoutis et al. 2002; Padioleau and Ridoux 2003; Nugent et al. 2003; Wright et al. 2003; Fraser and Chang 2003; Tolia et al. 2003; Papathanasiou and Scott 2004; Prabhakaran et al. 2005; Tan et al. 2005; Muniswamy-Reddy et al. 2006; Shrira and Xu 2006].

3. SUGGESTED BENCHMARKING GUIDELINES

We now present a list of guidelines to consider when evaluating the performance of a file or storage system. A Web version summary of this document can be found at `www.fsl.cs.sunysb.edu/project-fsbench.html`.

The two underlying themes are:

(1) Explain *exactly* what you did:
For example, if you decided to create your own benchmark, describe it in detail. If you are replaying traces, describe where they are from, how they were captured, and how you are replaying them (what tool? what speed?). This can help others understand and validate your results.

(2) Do not just say *what* you did, but justify *why* you did it that way:
For example, while it is important to note that you are using ext2 as a baseline for your analysis, it is just as important (or perhaps even *more* important) to discuss why it is a fair comparison. Similarly, it is useful for the reader to know *why* you ran that random-read benchmark so that they know what conclusions to draw from the results.

3.1 Choosing The Benchmark Configurations

The first step of evaluating a system is to pose questions that will reveal the performance characteristics of the system, such as "how does my system compare to current similar systems?," "how does my system behave under its expected workload?," and "what are the causes of my performance improvements or overheads?" Once these questions are formulated, one must decide on what baseline systems, system configurations, and benchmarks should be used to best answer them. This will produce a set of $\langle$*system*, *configuration*, *benchmark*$\rangle$ tuples that will need to be run. The researcher should have a rough idea what the results should be for each configuration at this point; if the actual results differ from these expectations, then the causes of the deviations must be investigated.

Since a system's performance is generally more meaningful when compared to the performance of existing technology, one should find existing systems that provide fair and interesting comparisons. For example, for benchmarking an encryption storage device, it would be useful to compare the performance to other encrypted storage devices, a traditional device, and perhaps some alternate implementations (user-space, file system, etc.).

The system under test may have several configurations that will need to be evaluated in turn. In addition, one may create artificial configurations where a component of the system is removed to determine its overhead. For example, in an encryption file or storage system, you can use a null cipher (copy data only) rather than encrypt, to isolate the overhead of encryption. Determining the cause of overheads may also be done using profiling techniques. Showing a breakdown of performance numbers helps one to fully understand a system's behavior and is generally a good practice.

There are three main types of benchmarks that one can choose from: macro-benchmarks, trace replaying, and micro-benchmarks.

Macro-benchmarks. These exercise multiple file system operations, and are usually good for an overall view of the system's performance, though the workload may not be realistic. These benchmarks are described further in Section 7.

Trace-based. Replaying traces can also provide an overall view of the system's performance. Traces are usually meant to exercise the system with a representative real-world workload, which can help to better understand how a system would behave under normal use. However, one must ensure that the trace is in fact representative of that workload (for example, the trace should capture a large enough sample), and that the method used to replay the trace preserves the characteristics of the workload. Section 8 provides more information about trace-based benchmarking.

Micro-benchmarks. These exercise few (usually one or two) operations. These are useful if you are measuring a very small change, to better understand the results of a macro-benchmark, to isolate the effects of specific parts of the system, or to show worst-case behavior. In general, these benchmarks are only meaningful when presented together with other benchmarks. See Section 9 for more information.

We recommend using at least one macro-benchmark or trace to show a high-level view of performance, along with several micro-benchmarks to highlight more focused views. In addition, there are several workload properties that should be considered. We describe five important ones here. First, benchmarks may be characterized by how CPU or I/O bound they are. File and storage system benchmarks should generally be I/O-bound, but a CPU-bound benchmark should also be run for systems that exercise the CPU. Second, if the benchmark records its own timings, it should use accurate measurements. Third, the benchmark should be scalable, meaning that it exercises each machine the same amount, independent of hardware or software speed. Fourth, multi-threaded workloads may provide more realistic scenarios, and may help saturate the system with requests. Fifth, the workloads should be well-understood. While the code of synthetic benchmarks can be read, and traces can be analyzed, it is more difficult to understand some application workloads. For example, compile benchmarks can behave very differently depending on the testbed's architecture, installed software, and the version of the software being compiled. The source code for ad-hoc benchmarks should be publicly released, as it is the only truly complete description of your benchmark that would allow others to reproduce it (including any bugs or unexpected behavior).

3.2 Choosing The Benchmarking Environment

The state of the system during the benchmark's runs can have a large effect on results. After determining an appropriate state, it should be created accurately and reported along with the results. Some major factors that can affect results are cache state, ZCAV effects, file system aging, and non-essential processes running during the benchmark.

The state of the system's caches can affect the code-paths that are tested and thus affect benchmark results. It is not always clear if benchmarks should be run with "warm" or "cold" caches. On one hand, real systems do not generally run with completely cold caches. On the other hand, a benchmark that accesses too much cached data may be unrealistic as well. In addition, since requests will be mainly serviced from memory, the file or storage system will not be adequately exercised. Further, not bringing the cache back to a consistent state between runs can cause timing inconsistencies. If cold cache results are desired, caches should be cleared before each run. This can be done by allocating and freeing large amounts of memory, remounting the file system, reloading the storage driver, or rebooting. However, we have found that rebooting is more effective than the other methods [Wright et al. 2005]. When working in an environment with multiple machines, the caches on all necessary machines must be cleared. This will help create identical runs, thus ensuring more stable results. If, however, warm cache results are desired, this can be achieved by running the experiment n+1 times, and discarding the first run's result.

Most modern disks use Zoned Constant Angular Velocity (ZCAV) to store data. In this design, the cylinders are divided into zones, where the number of sectors in a cylinder increases with the distance from the center of the disk. Because of this, the transfer rate varies from zone to zone [Meter 1997]. It has been recommended to minimize ZCAV effects by creating a partition of the smallest possible size on the outside of the disk [Ellard and Seltzer 2003b]. However, this makes results less realistic, and may not be appropriate for all benchmarks (for example, long seeks may be necessary to show the effectiveness of the system). We recommend simply specifying the location of the test partition in the paper to help reproducibility.

Most file system and storage benchmarks are run on an empty system, which could make the results different than a real-world setting. A system may be aged by running a workload based on system snapshots [Smith and Seltzer 1997]. However, aging a 1GB file system by seven months using this method required writing 87.3GB of data. The amount of time required to age a file system would make it impractical for larger systems. TBBT has a faster, configurable aging technique, but it is less realistic since it is purely synthetic [Zhu et al. 2005]. Some other ways to age a system before running a benchmark are to run a long-term workload, copy an existing raw image, or to replay a trace before running the benchmark. It should be noted that for some systems and benchmarks, aging is not a concern. For example, aging will not have any effect when replaying a block-level trace on a traditional storage device, since the benchmark will behave identically regardless of the disk's contents.

To ensure the reproducibility of results, all non-essential services and processes should be stopped before running the benchmark. These processes can cause anomalous results or higher than normal standard deviations for a set of runs. However, processes such as `cron` will coexist with the system when used in the real world, and so it must be understood that these results are measured in a sterile environment. Ideally, we would be able to demonstrate performance with the interactions of other processes present. However, this is difficult because the set of processes is specific to a machine's configuration. Instead, we recommend using multi-threaded workloads because they more accurately depict a real system that normally has several active processes. In addition, ensure that no users log into the test machines, and make sure that no other traffic is consuming your network bandwidth while running benchmarks that involve your network.

3.3 Running The Benchmarks

There are four important guidelines to running benchmarks properly. First, one should ensure that every benchmark run is identical. Second, each test should be run several times to ensure accuracy, and standard deviation or confidence levels should be used to determine the appropriate number of runs. Third, tests should be run for a long enough period of time, so that the system is in steady state for a majority of the run. Fourth, the benchmarking process should be automated using scripts or available tools such as Auto-pilot [Wright et al. 2005] to avoid mistakes associated with manual repetitive tasks. This is discussed further in Section 11.

3.4 Presenting The Results

Once results are obtained, they must be presented appropriately so that accurate conclusions may be derived from them. Aside from the data that is presented, the benchmark configurations and environment should be accurately described. Proper graphs should be displayed, with error bars, where applicable.

We recommend using confidence intervals, rather than standard deviation, to present results. The standard deviation is a measure of how much variation there is between the runs. The half-width of the confidence interval describes how far the true value may be from the captured mean with a given degree of confidence (e.g., 95%). This provides a better sense of the true mean. In addition, as more benchmark runs are performed, the standard deviation may not decrease, but the width of the confidence interval will.

For experiments with less than 30 runs, one should be careful not to use the normal distribution for calculating confidence intervals. This is because the central limit theorem no longer holds with a small sample size. Instead, one must use the Student's t-distribution. This distribution may also be used for experiments with at least 30 runs, since in this case it is similar to the normal distribution.

Large confidence interval widths or non-normal distributions may indicate a software bug or benchmarking error. For example, the half-widths of the confidence intervals should be less than 5% of the mean. If the results are not stable, then either there is a bug in the code, or the instability should be explained. Anomalous results (e.g., outliers) should never be discarded. If they are due to programming or benchmarking errors, the problem should be fixed and the benchmarks rerun to gather new and more stable results.

3.5 Validating Results

Other researchers may wish to benchmark your software for two main reasons: (1) If they wish to reproduce your results or confirm them, or (2) If they want to compare their system to yours.

First, it is considered good scientific practice to provide enough information for others to validate your results. This includes detailed hardware and software specifications about the testbeds. Although it is usually not practical to include such large amounts of information in a conference paper, it can be published in an online appendix. Whereas it can be difficult for a researcher to accurately validate another's results without the exact testbed, it is still possible to see if the results generally make sense.

Second, there may be a case where a researcher creates a system that has similar properties to yours (e.g., they are both encryption file systems). It would be logical for the researcher to compare the two systems. However, if your paper showed an X% overhead over ext2, and the new file system has a Y% overhead over ext2, no claim can be made about which of the two file systems is better. The researcher should benchmark both file systems, using a setup that is as similar as possible to that of the original benchmark. This way both file systems are tested under the same conditions. Moreover, since they are running the benchmark in the same way that you did, no claim can be made that they chose a specific case in which their file system performs better.

To help solve these two issues, enough information must be made available about your testbed (both hardware and any relevant software) so that an outside researcher can validate your results. If possible, make your software available to other researchers so that they can compare their system to yours. Releasing the source is preferred, but a binary release can also be helpful if there are legal issues preventing the release of source code. SOSP 2007 is attempting to improve this situation by asking authors in the submission form if they will make the source code and raw data for their system and experiments available so that others can reproduce the results. If enough authors agree to this sharing, and other conferences follow suit, this could make it easier to compare similar systems and reproduce results in the future. In addition, any benchmarks that were written and any traces that were collected should be made available to others.

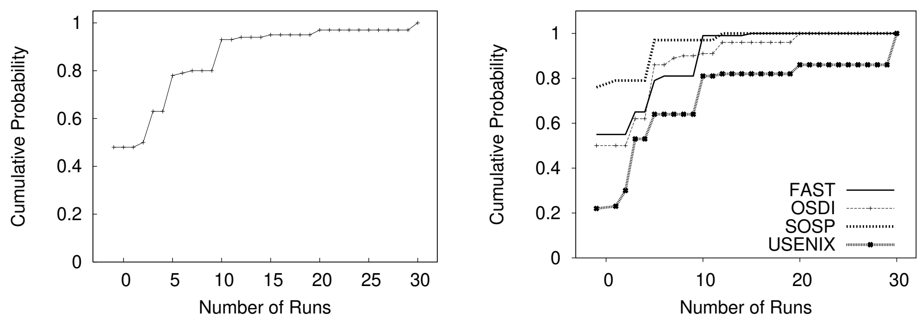

(a) Number of runs of surveyed benchmarks.

(b) Number of runs categorized by conference.

Fig. 1. The figure on the left shows the cumulative distribution function (CDF) for the number of runs that were performed in the surveyed benchmarks. The figure on the right is the CDF of the same data categorized by conference. A value of –1 was used for benchmarks where the number of runs was not specified.

4. COMPLIANCE WITH THE GUIDELINES

We now examine how well the surveyed papers followed the benchmarking practices that were discussed in Section 3.

Number of runs. Running benchmarks multiple times is important for ensuring accuracy and presenting the range of possible results. Reporting the number of runs allows the reader to determine the benchmarking rigor. We now examine the number of runs performed in each surveyed experiment. To ensure accuracy, we did not include experiments where one operation was executed many times and the per-operation latency was reported, because it was not clear whether to count the number of runs as the number of times the operation was executed, or the number of times the entire benchmark was run. Figure 1 shows the results from the 388 benchmarks that were counted. We found that two papers [Wang et al. 2002; Wright et al. 2003] ran their benchmarks more than once, since they included error bars or confidence intervals, but did not specify the number of runs. These are shown as two runs. The figure shows that the number of runs were not specified for the majority of benchmarks. Assuming that papers that did not specify the number of runs ran their experiments once, we can break down the data by conference:

Conference	Mean	Standard Deviation	Median
SOSP	2.1	2.4	1
FAST	3.6	3.6	1
OSDI	3.8	4.3	2
USENIX	4.7	6.2	3

The per-conference values are presented for informational value and we feel they may of interest to the reader. However, we caution the reader against drawing conclusions based on these statistics, as benchmarking rigor alone does not determine the quality of a conference, and the number of runs alone does not determine benchmarking rigor.

Statistical dispersion. After performing a certain number of runs, it is important to inform the reader about the statistical dispersion of the results. 34.6% of the surveyed papers included at least a general discussion of standard deviation, and 11.2% included confidence

Table I. Percentage of papers that discussed standard deviations or confidence intervals, classified by conference.

	FAST	OSDI	SOSP	USENIX
Number of papers	12	21	51	23
Standard deviations	8.3%	28.6%	27.5%	69.6%
Confidence intervals	16.7%	19.1%	7.8%	8.7%
Total	25.0%	47.6%	35.3%	78.3%

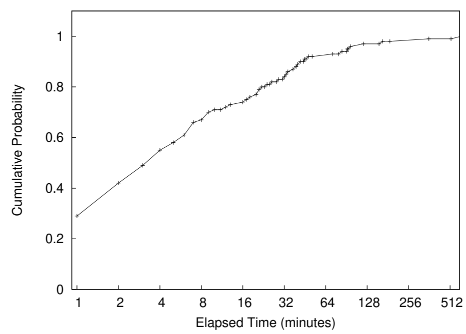

Fig. 2. CDF of the number of benchmarks that were run in the surveyed papers with a given elapsed time. Note the log scale on the x-axis.

intervals. The percentage of papers that have discussed either one varied between 35.7% and 83.3% per year, but there was no upward or downward trend over time. Interestingly, we did notice significant differences between conferences, shown in Table I, but we do not suggest that this is telling of the overall quality of any particular conference. In addition to informing the reader about the overall deviations or intervals for the paper, it is important to show statistical dispersion for each result. This can be done with error bars in graphs, by augmenting tables, or by mentioning it in the text. From all of the surveyed benchmarks, only 21.5% included this information.

Benchmark runtimes. To achieve stable results, benchmarks must run for a long enough time to reach steady state and exercise the system. This is especially important as benchmarks must scale with increasingly faster hardware. We looked at the runtimes of the 198 experiments that specified the elapsed time of the benchmark. Most benchmarks that reported only per-operation latency or throughput did not specify their runtime. For each experiment, we took the longest elapsed time of all configurations, and rounded them up to the nearest minute. For benchmarks with multiple phases, times were added to create a total time. The results are summarized in Figure 2. We can see in the figure that 28.6% of benchmarks ran for less than one minute, 58.3% ran for less than five minutes, and 70.9% ran for less than ten.

Number of benchmarks. The number of benchmarks used for performance evaluations in each paper is shown in Figure 3. We can see that 37.7% of the papers used only one or two benchmarks, which in most cases is not sufficient for a reader to fully understand the performance of a system.

System descriptions. To gain some idea of testbed specifications that were published in the surveyed papers, we now present the number of parameters that were listed. It must be noted that not all parameters are equally important, and that some parameters are actually

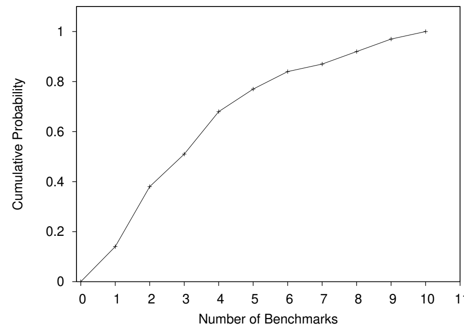

Fig. 3. CDF of the number of benchmarks that were run in the surveyed papers.

a subset of other parameters. For example, a disk's speed is counted as one parameter, but a disk's model number is counted as one parameter as well, even though the disk's speed and several other disk parameters can be found from the model specifications. Since it is not clear how to weigh each parameter, we will instead caution that these results should be used only as rough estimates. An average of 7.3 system parameters were reported per paper, with a standard deviation of 3.3. The median was 7. While this is not a small number, it is not sufficient for reproducing results. In addition, only 35.9% of the papers specified the cache state during the benchmark runs. We specify the testbed used in this article in Section 6, which we believe should be sufficient to reproduce our results.

5. RELATED WORK

A similar survey was conducted in 1997 covering more general systems papers [Small et al. 1997]. The survey included ten conference proceedings from the early to mid 90's. The main goals of that survey were to determine how reproducible and comparable the benchmarks were, as well as to discuss statistical rigor. We do not discuss statistical rigor in more detail in this paper, since there is a good discussion presented there. They went on to advise on how to build good benchmarks and report them with proper statistical rigor. Some results from their survey are that over 90% of file system benchmarks run were ad-hoc, and two-thirds of experiments presented a single number as a result without any statistical information.

In 1999, Mogul presented a similar survey of general OS research papers and commented on the lack of standardization of benchmarks and metrics [Mogul 1999]. He conducted a small survey of two conference proceedings and came to the conclusion that the operating system community is in need of good standardized benchmarks. Of the eight file system papers he surveyed, no two used the same benchmark to analyze performance.

Chen and Patterson concentrated on developing an I/O benchmark that can shed light on the causes for the results, scales well, has results that are comparable across machines, is general enough that it can be used by a wide range of applications, and is tightly specified so that everyone follows the same rules [Chen and Patterson 1993]. This benchmark does not perform metadata operations, and is designed to benchmark drivers and I/O devices. The authors go on to discuss how they made a self-scaling benchmark with five parameters: data size, average size of an I/O request, fraction of read operations (the fraction of write operations is 1 minus this value), fraction of sequential access (the fraction of random access is 1 minus this value), and the number of processes issuing I/O requests. Their benchmark keeps four of these parameters constant while varying the fifth, producing five graphs. Because self-scaling will produce different workloads on different machines, the paper discusses how to predict performance, so that results can be compared, and shows reasonable ability to perform the predictions.

A paper by Tang and Seltzer from the same research group, entitled "Lies, Damned Lies, and File System Benchmarks" [Tang and Seltzer 1994], was one source for some of our observations about the Andrew (Section 7.3), LADDIS (Section 7.5), and Bonnie (Section 9.1) benchmarks, as well as an inspiration for this larger study.

Tang later expanded on the ideas of this paper, and introduced a benchmark called `dtangbm` [Tang 1995]. This benchmark consists of a suite of micro-benchmarks called *fsbench* and a workload characterizer. Fsbench has four phases:

(1) Measures the disk performance so that it can be known whether improvements are due to the disk or file system.

(2) Estimates the size of the buffer cache, the attribute cache, and the name translation cache. This information is used by the next two phases to ensure proper benchmark scaling.

(3) Runs the micro-benchmarks whose results are reported. The benchmark takes various measurements within each micro-benchmark, providing much information about the file system's behavior. The reported metric is KB/sec. The first two micro-benchmarks in this phase test block allocation to a single file for sequential and random-access patterns. The third micro-benchmark tests how blocks are allocated to files that are in the same directory. The fourth micro-benchmark measures the performance of common meta-data operations (create, delete, `mkdir`, `rmdir`, and `stat`).

(4) Performs several tests to help file system designers pinpoint performance problems. It isolates latencies for attribute (inode) creation, directory creation, attribute accesses, and name lookups by timing different meta-data operations and performing some calculations on the results. It also uses a variety of read patterns to find cases where read-ahead harms performance. Finally, it tests how well the file system handles concurrent requests.

The second component of `dtangbm`, the workload characterizer, takes a trace as input, and prints statistics about the operation mix, sequential versus random accesses, and the average number of open files. This information could theoretically be used in conjunction with the output from `fsbench` to estimate the file system's performance for any workload, although the authors of `dtangbm` were not able to do so accurately in that work.

Another paper from Seltzer's group [Seltzer et al. 1999] suggests that not only are currently-used benchmarks poor, but the types of benchmarks that are run do not provide much useful information. The current metrics do not provide a clear answer of which system would perform better for a given workload. The common and simple workloads are not adequate, and so they discuss three approaches to application-specific benchmarking. In the first, system properties are represented in one vector and the workload properties are placed in another. Combining the two vectors can produce a relevant performance metric. The second approach involves using traces to develop profiles that can stochastically generate similar loads. The third uses a combination of the first two.

According to Ruwart, not only are current benchmarks ill-suited for testing today's systems, they will fare even worse in the future because of the systems' growing complexities (e.g., clustered, distributed, and shared file systems) [Ruwart 2001]. He discusses an approach to measuring file system performance in a large-scale, clustered supercomputer environment while describing why current techniques are insufficient.

Finally, Ellard and Seltzer describe some problems that they experienced while benchmarking a change to an NFS server [Ellard and Seltzer 2003b]. First, they describe ZCAV effects, which were previously documented only in papers that discuss file system layouts [Meter 1997] (not in performance evaluations). Since the inner tracks on a disk have fewer sectors than the outer tracks, the amount of data read in a single revolution can vary greatly. Most papers do not deal with this property. Aside from ZCAV effects, they also describe other factors that can affect performance, such as SCSI command queuing, disk scheduling algorithms, and differences between transport protocols (i.e., TCP and UDP).

6. BENCHMARKING METHODOLOGY

In this section, we present the testbed and benchmarking procedures that we used for conducting the experiments throughout the remainder of this paper. We describe the hardware and software configuration of the test machine in Section 6, and we discuss our benchmarking procedure in Section 6.

System Configuration. We conducted all our experiments on a machine with a 1.7GHz Pentium 4 CPU, 8KB of L1 Cache, and 256KB of L2 Cache. The motherboard was an Intel Desktop Board D850GB with a 400 MHz System Bus. The machine contained 1GB of PC800 RAM. The system disk was a 7200 RPM WD Caviar (WD200BB) with 20GB capacity. The benchmark disk was a Maxtor Atlas (Maxtor-8C018J0) 15,000 RPM, 18.4GB, Ultra320 SCSI disk. The SCSI controller was an Adaptec AIC-7892A U160.

The operating system was Fedora Core 6, with patches as of March 07, 2007. The system was running a vanilla 2.6.20 kernel and the file system was ext2, unless otherwise specified. Some relevant program versions, obtained by passing the `--version` flag on the command line, along with the Fedora Core package and version are GCC 4.1.1 (gcc.i386 4.1.1-51.fc6), GNU ld 2.17.50.0.6-2.fc6 (binutils 2.17.50.0.6-2.fc6), GNU autoconf 2.59 (autoconf.noarch 2.59-12), GNU automake 1.9.6 (automake.noarch 1.9.6-2.1), GNU Make 3.81 (make 1:3.81-1.1), and GNU tar 1.15.1 (tar 2:1.15.1-24.fc6).

The kernel configuration file and the full package listing are available at `www.fsl.cs.sunysb.edu/project-fsbench.html`.

Benchmarking Procedure. We used the Autopilot v.2.0 [Wright et al. 2005] benchmarking suite to automate the benchmarking procedure. We configured Autopilot to run all tests at least ten times, and compute 95% confidence intervals for the mean elapsed, system, and user times using the Student-*t* distribution. In each case, the half-width of the interval was less than 5% of the mean. We report the mean of each set of runs. In addition, we define "wait time" to be the time that the process was not using the CPU (mostly due to I/O).

Autopilot rebooted the test machine before each new sequence of runs to minimize the influence of different experiments to each other. Autopilot automatically disabled all unrelated system services to prevent them from influencing the results. Compilers and executables were located on the machine's system disk, so the first run of each set of tests was discarded to ensure that the cache states were consistent. We configured Autopilot to unmount, recreate and then remount all tested file systems before each benchmark run. To minimize ZCAV effects, all benchmarks were run on a partition located toward the outside of the disk that was just large enough to accommodate the test data [Ellard and Seltzer

2003b]. However, the partition size was big enough to avoid the file system's space-saving mode of file system operation. In the space-saving mode, file systems optimize their operation to save disk space and thus have different performance characteristics [Meter 1997].

7. MACRO-BENCHMARKS

In this section we describe the macro-, or general-purpose, benchmarks that were used in the surveyed research papers. We point out the strengths and weaknesses in each. For completeness, we also discuss several benchmarks that were not used. Macro-benchmark workloads consist of a variety of operations and aim to simulate some real-world workload. The disadvantage of macro-benchmarks is that the workload may not be representative of the workload that the reader is interested in, and it is very difficult to extrapolate from the performance of one macro-benchmark to a different workload.

Additionally, there is no agreed-upon file system benchmark that everyone can use. Some computer science fields have organizations that create benchmarks and keep them up to date (TPC in the database community, for example). There is no such organization specifically for the file system community, although SPEC, has one benchmark targeted for a specific network file system protocol—see Section 7.5. For storage, the Storage Performance Council [SPC 2007] has created two standardized benchmarks, which we describe in Section 7.6. We have observed that many researchers use the same benchmarks, but they often do not explain the reasons for using them or what the benchmarks show about the systems they are testing. From the 148 macro-benchmark experiments that were performed in the surveyed papers, 20 reported that they have done so because the benchmark was popular or standard, and 28 provided no reason at all. Others described what real-world workload the given benchmark was mimicking, but did not say why it was important to show those results. In total, inadequate reasoning was given for at least 32.4% of the macro-benchmark experiments that were performed. This leads us to believe that many researchers use the benchmarks that they are used to and that are commonly used, regardless of suitability.

We describe Postmark in Section 7.1, compile benchmarks in Section 7.2, the Andrew benchmark in Section 7.3, TPC benchmarks in Section 7.4, SPEC benchmarks in Section 7.5, SPC benchmarks in Section 7.6, NetNews in Section 7.7, and other macro-benchmarks in Section 7.8.

7.1 Postmark

Postmark [Katcher 1997; VERITAS Software 1999], created in 1997, is a single-threaded synthetic benchmark aimed at measuring file system performance over a workload composed of many short-lived, relatively small files. Such a workload is typical of electronic mail, Netnews, and Web-based commerce transactions as seen by ISPs. The workload includes a mix of data and meta-data–intensive operations. However, the benchmark only approximates file system activity: it does not perform any application processing, and so the CPU utilization is less than that of an actual application.

The benchmark begins by creating a pool of random text files with uniformly distributed sizes within a specified range. After creating the files, a sequence of "transactions" is performed (in this context, a transaction is a Postmark term, and is unrelated to the database concept). The number of files, the number of subdirectories, the file size range, and the number of transactions are all configurable. Each Postmark transaction has two parts: a file creation or deletion operation paired with a file read or append. The ratios of reads to

Table II. The default Postmark v1.5 configuration, along with the number of research papers that disclosed each piece of information (from the 30 papers that used Postmark in the papers we surveyed).

Parameter	Default Value	Number Disclosed (out of 30)
File sizes	500–10,000 bytes	21
Number of files	500	28
Number of transactions	500	25
Number of subdirectories	0	11
Read/write block size	512 bytes	7
Operation ratios	equal	16
Buffered I/O	yes	6
Postmark Version	-	7

appends and creates to deletes is configurable. A file creation operation creates and writes random text to a file. A file deletion operation removes a randomly chosen file from the active set. A file read operation reads a random file in its entirety and a file write operation appends a random amount of data to a randomly chosen file. It is also possible to choose whether or not to use buffered I/O.

One drawback of using Postmark is that it does not scale well with the workload. Its default workload, shown in Table II, does not exercise the file system enough. This makes it no longer relevant to today's systems, and as a result researchers use their own configurations. On the machine described in Section 6, the default Postmark configuration takes less than a tenth of one second to run, and barely performs any I/O. One paper [Nightingale et al. 2005] used the default configuration over NFS rather than updating it for current hardware, and the benchmark completed in under seven seconds. It is unlikely that any accurate results can be gathered from such short benchmark runs. In Section 12.2, we show how other Postmark configurations behave very differently from each other. Rather than having the number of transactions to be performed as a parameter, it would be more beneficial to run for a specified amount of time, and report the peak transaction rate achieved. Benchmarks such as Spec SFS and AIM7 employ a similar methodology.

Having default parameters that become outdated creates two problems. First, there is no standard configuration, and since different workloads exercise the system differently, the results across research papers are not comparable. Second, not all research papers precisely describe the parameters that were used, and so results are not reproducible.

Few research papers specify all parameters necessary for reproducing a Postmark benchmark. From the 106 research papers that we surveyed, 30 used Postmark as one of their methods for performance evaluation [Ng et al. 2002; Sarkar et al. 2003; Sivathanu et al. 2004; Radkov et al. 2004; Tan et al. 2005; Anderson et al. 2002; Nightingale et al. 2005; Wang et al. 2002; Seltzer et al. 2000; MacCormick et al. 2004; Abd-El-Malek et al. 2005; Thereska et al. 2004; Magoutis et al. 2002; Aranya et al. 2004; Muniswamy-Reddy et al. 2004; Zhang and Ghose 2003; Wright et al. 2003; Strunk et al. 2000; Nightingale et al. 2006; Sivathanu et al. 2004; Prabhakaran et al. 2005; Stein et al. 2001; Soules et al. 2003; Sivathanu et al. 2006; Denehy et al. 2005; Weddle et al. 2007; Schindler et al. 2002; Magoutis et al. 2003; Wang et al. 2002; Wang and Merchant 2007]. Table II shows how many of these papers disclosed each piece of information. Papers that use configurable benchmarks should include all parameters to make the results meaningful; only five did so. These five papers specified that any parameters not mentioned were the defaults (Table II gives them credit for specifying all parameters).

In addition to failing to specify parameters, Table II shows that only 5 out of 30 research papers mentioned the version of Postmark that they used. This is especially crucial with Postmark because of major revisions that make results from different versions incomparable. The biggest changes were made in version 1.5, where the benchmark's pseudo-random number generator was overhauled. Having a generator in the program itself is a good idea, as it makes benchmarks across various platforms more comparable. There were two key bugs with the previous pseudo-random number generator. First, it did not provide numbers that were random enough. Second, and more importantly, it did not generate large enough numbers, so files were not created as large as the parameter specified, causing results to be inaccurate at best. Having a built-in pseudo-random number generator is an example of a more general rule: library routines should be avoided unless the goal of the benchmark is to measure the libraries because this introduces more dependencies on the machine setup (OS, architecture, and libraries).

Another lesson that Postmark teaches us is to make an effort to keep benchmarking algorithms scalable. The algorithm that Postmark uses to randomly choose files is $O(N)$ on the number of files, which does not scale well with the workload. It would be trivial to modify Postmark to fix this, but would make results incomparable with others. While high levels of computation are not necessarily a bad quality, they should be avoided for benchmarks that are meant to be I/O-bound.

An essential feature for a benchmark is accurate timing. Postmark uses the `time(2)` system call internally, which has a granularity of one second. There are better timing functions available (e.g.,`gettimeofday`) that have much finer granularity and therefore provide more meaningful and accurate results.

One of the future directions that Postmark was looking at is allowing different numbers of readers and writers, instead of just one process that does both. Four of the surveyed papers [Aranya et al. 2004; Anderson et al. 2004; Anderson et al. 2002; Wang and Merchant 2007] ran concurrent Postmark processes. Seeing how multiple processes affect results is useful for benchmarking most file systems, as this reflects real-world workloads more closely. However, since Postmark is not being maintained (no updates have been made to Postmark since 2001), this will probably not be done.

One research paper introduces Filemark [Bryant et al. 2002], which is a modified version of Postmark 1.5. It differs in five respects. First, it adds multi-threading so that it can produce a heavier and more realistic workload. Second, it uses `gettimeofday` instead of `time` so that timing is more accurate. Third, it uses the same set of files for multiple transaction phases. This makes the runtime faster, but performs fewer writes, and extra care must be taken to ensure that data is not cached if this is not desired. Fourth, it allows the read-write and create-delete ratios to be specified to the nearest 1% instead of 10% as with Postmark. Fifth, it adds an option to not perform the delete phase, which the Filemark authors claim has a high variation and is almost meaningless. We suggest instead that if some operation has a high variation, it should be further investigated and explained rather than discarded.

Postmark puts file systems under heavy stress when the configuration is large enough, and is a fairly good benchmark. It has good qualities such as a built in pseudo-random number generator, but also has some deficiencies. The important thing is to keep the positive and negative qualities in mind when running the benchmark and analyzing results. In sum, we suggest that Postmark be improved to have a scalable workload, more accurate timing, and allow for multi-threaded workloads.

7.2 Compile Benchmarks

Thirty-six of the papers that we surveyed timed the compiling of some code to benchmark their projects:

—ten compiled SSH [Soules et al. 2003; Ng et al. 2002; Schindler et al. 2002; Denehy et al. 2005; Strunk et al. 2000; Prabhakaran et al. 2005; Seltzer et al. 2000; Kroeger and Long 2001; Stein et al. 2001; Weil et al. 2006],

—eight compiled an OS kernel [Zhang and Ghose 2003; Mazieres et al. 1999; Tolia et al. 2004; Gulati et al. 2007; Radkov et al. 2004; Sivathanu et al. 2006; Papathanasiou and Scott 2004; Muniswamy-Reddy et al. 2006],

—six papers (all from our research group) compiled Am-utils [Aranya et al. 2004; Muniswamy-Reddy et al. 2004; Zadok and Nieh 2000; Zadok et al. 2001; Wright et al. 2003; Zadok et al. 1999],

—five compiled Emacs [Gulati et al. 2007; Li et al. 2004; Fu et al. 2000; Maziéres 2001; Muthitacharoen et al. 2001],

—three compiled Apache [Peek and Flinn 2006; Nightingale et al. 2006; Nightingale et al. 2005],

—five compiled other packages [Gulati et al. 2007; Sobti et al. 2002; Abd-El-Malek et al. 2005; Gniady et al. 2004; Lee et al. 1999],

—two did not specify the source code being compiled [Kroeger and Long 2001; Kim et al. 2000].

The main problem with compile benchmarks is that because they are CPU-intensive, they can hide overheads in many file systems. This issue is discussed further in Section 12.2. However, a CPU-intensive benchmark may be a reasonable choice for a file system that already has a significant CPU component (such as an encryption or compression file system). Even so, a more I/O-intensive benchmark should be run as well. Other issues relating to compile benchmarks affect the ability of readers to compare, fully understand, and reproduce benchmark results:

—Different machines may have different compiler tool chains, specifically:
  —Different architectures produce different code.
  —Different compilers utilize different optimizations.
  —The source code may not compile on all architectures, and older packages often cannot compile on newer systems so the workload becomes obsolete.

—Different machines are configured differently (with regard to both hardware and software), so the configuration phase will not be the same on all machines, and the resulting code will be different as well.

—Some compilations (such as kernels) have different configuration options, resulting in a different configuration phase, and different resulting code. Although default configuration files are sometimes included, using them can result in compilation errors, as we have experienced when compiling some versions of the Linux kernel.

—The operation mixes change depending on which program is being compiled, and even on its version.

—The compile process is always growing more complex and there is much variation between programs. Most require explicit configuration phases, and some require phases

Table III. SSH 2.1.0, Am-utils 6.1b3, and Linux kernel 2.4.20 characteristics. The total size refers to the pre-compiled package, since the total size after compiling is system-dependent.

	SSH	Am-utils	Linux Kernel
Directories	54	25	608
Files	637	430	11,352
Lines of Code	170,239	61,513	4,490,349
Code Size (Bytes)	5,313,257	1,691,153	126,735,431
Total Size (Bytes)	9,068,544	8,441,856	174,755,840

that resolve dependencies. The amount of time spent in each phase can also vary significantly [Zadok 2002].

To allow a benchmark to be accurately reproduced, all parameters that could affect the benchmark must be reported. This is particularly difficult with a compile benchmark. From the 33 papers that used compile benchmarks, only one specified the compiler and linker versions, and one specified compiler options. Eight failed to specify the version of the code that was being compiled, and 19 failed to specify the compilation steps that were being measured. Although it is easy to report the source code version, it is more difficult to specify relevant programs and patches that were installed. For example, Emacs has dependencies on the graphical environment, which may include dozens of libraries and their associated header files. However, this is feasible to do if a package manager has been used and can provide information about all of the installed program versions. Because of the amount of information that needs to be presented, we recommend creating an online appendix with the detailed testbed setup.

There is a common belief that file systems see similar loads independent of the software being compiled. Using OSprof [Joukov et al. 2006], we profiled the build process of three packages commonly used as compile benchmarks: (1) SSH 2.1.0, (2) Am-utils 6.1b3, and (3) the Linux 2.4.20 kernel with the default configuration. Table III shows the general characteristics of the packages. The build process of these packages consists of a configuration and a compilation phase. The configuration phase consists of running GNU `configure` scripts for SSH and Am-utils, and running "`make defconfig dep`" for the Linux kernel. We analyzed the configuration and compilation phases separately, as well as together. Before the configuration and compilation phases, we remounted the ext2 file system that the benchmark was run on to reduce caching effects. Figure 4 shows the distribution of the total number of invocations of all the ext2 VFS operations used during the build processes. Note that each of the three graphs uses different scales for the number of operations (y-axis).

Figures 4(a) and 4(b) show that even though the SSH and Am-utils build process sequence, source-file structure, and total sizes appear to be similar, their operation mixes are quite different; moreover, the fact that SSH has nearly three times the lines of code of Am-utils is also not apparent from analyzing the figures. In particular, the configuration phase dominates in the case of Am-utils whereas the compilation phase dominates the SSH build. More importantly, the read-write ratio for the Am-utils build was 0.75:1, whereas it was 1.28:1 for the SSH build. This can result in significant performance differences for read-oriented or write-oriented systems. Not surprisingly, the kernel build process's profile differs from both SSH and Am-utils. As can be seen in Figure 4(c), both of the kernel build phases are strongly read biased. In addition, the kernel build process is more intensive in file open and file release operations. As we can see, even seemingly similar

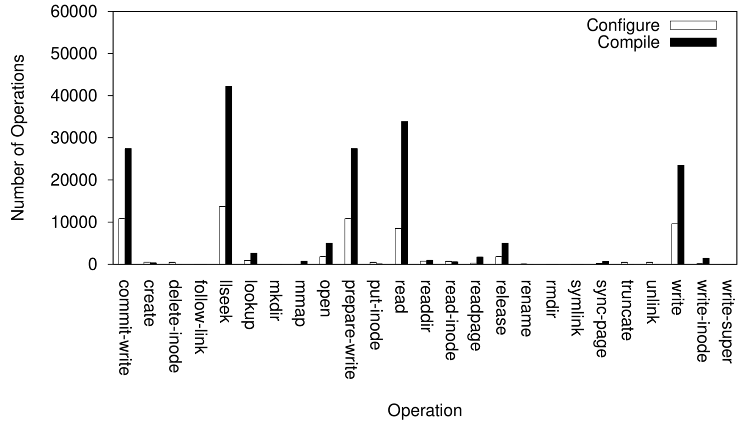

(a) SSH 2.1.0

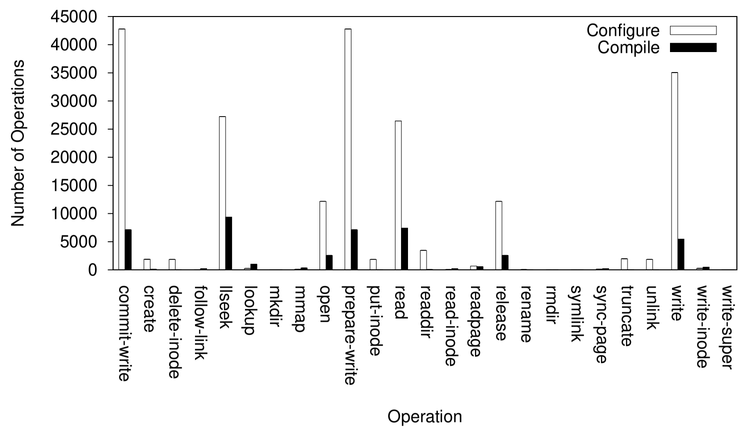

(b) Am-utils 6.1b3

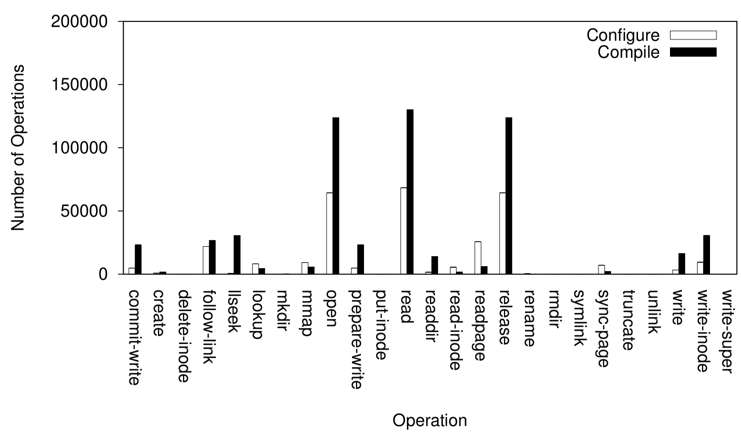

(c) Kernel 2.4.20

Fig. 4. Operation mixes during three compile benchmarks as seen by the ext2 file system. Note that each plot uses a different scale on the y-axis.

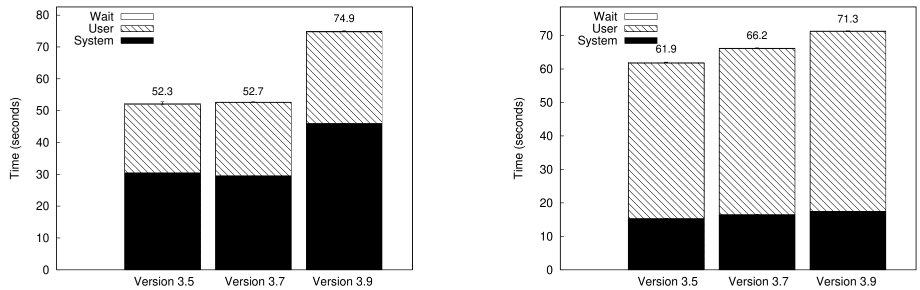

(a) Configure phase

(b) Compile phase

Fig. 5. Time taken to configure and compile OpenSSH versions 3.5, 3.7, and 3.9 on ext2. Note that error bars are shown, but are small and difficult to see.

compile benchmarks exercise the test file systems with largely different operation mixes.

Now let us consider compilation of the same software with slightly different versions. In "Opportunistic Use of Content Addressable Storage for Distributed File Systems," by Tolia, et al. [Tolia et al. 2003], the authors show the commonality found between versions of the Linux 2.4 kernel source code (from 2.4.0 to 2.4.20) and between several nightly snapshots of Mozilla binaries from March 16th, 2003 to March 25th, 2003. The commonality for both examples is measured as the percentage of identical blocks. The commonality between one version of the Linux source code and the next ranges from approximately 72% to almost 100%, and 2.4.20 has only about 26% in common with 2.4.0. The Mozilla binaries show us how much a normal user application can change over the course of one day—subsequent versions had approximately 42–71% in common, and only about 30% of the binary remained unchanged over the course of ten days. This illustrates the point that even when performing a compile benchmark on the same program, its version can greatly affect the results.

Not only do the source code and the resulting binaries change, but the operation mixes change as well. To illustrate this point, we compiled three different, but recent versions of SSH on our reference machine, using the same testbed and methodology described in Section 6. We used SSH because it is the most common application that was compiled in the papers we surveyed, and specifically OpenSSH because it compiles on modern systems.

Each test consisted of unpacking the source code, configuring it, compiling it, and then deleting it. The first and last steps are less relevant to our discussion, and so we do not discuss them further. The results for the configure and compile phases are shown in Figure 5. Although the elapsed times for the configure phase of versions 3.5 and 3.7 are indistinguishable, there is a much larger difference between versions 3.7 and 3.9 (42.3% more elapsed time, 55.6% more system time, and 25.1% more user time). There is a difference between all three versions for the compile phase, with increases ranging from 6.0% to 8.4% between subsequent versions for all time components. We can see that versions of the same program that are close to each other are very different, and we can therefore infer that the difference will be greater between versions that are spread further apart, and more so for different programs. Finally, we see how small the effects of I/O operations are on the benchmark results.

7.3 The Andrew File System Benchmark

This benchmark was created in 1988 to evaluate the performance of the Andrew File System [Howard et al. 1988]. The benchmark script operates on a directory subtree containing the source code for a program. The operations that were chosen for the benchmark were intended to be representative of an average user workload [Howard et al. 1988], although this was not shown to be statistically accurate. Nine papers that we surveyed used this benchmark for performance analysis [Adya et al. 2002; Grönvall et al. 1999; Rhea et al. 2003; Gopal and Manber 1999; Tolia et al. 2003; Kim et al. 2002; Aguilera et al. 2003; Saito et al. 2002].

The Andrew benchmark has five phases:

(1) *MakeDir* - Constructs directories in a target subtree identical to the structure of the original subtree.

(2) *Copy* - Copies all of the files from the source subtree to the target subtree.

(3) *ScanDir* - Perform a `stat` operation on each file of the target subtree.

(4) *ReadAll* - Read every byte of every file in the target subtree once.

(5) *Make* - Compile and link all files in the target subtree.

This benchmark has two major problems. First, the final phase of the benchmark (compilation) dominates the benchmark's run time, and introduces all of the drawbacks of compile benchmarks to this one (see Section 7.2). Second, the benchmark does not scale. The default data set will fit into the buffer cache of most systems today, so all read requests after the *Copy* phase are satisfied without going to disk. This does not provide an accurate picture of how the file system would behave under workloads where data is not cached. In order to resolve the issue of scalability, four of the research papers used a source program that is larger than the one that comes with the benchmark. This, however, causes results to be incomparable between papers.

Several research papers use a modified version of the Andrew benchmark (MAB) [Ousterhout 1990] from 1990. The modified benchmark uses the same compiler to make the results more comparable between machines. This solves one of the issues that was seen when examining compile benchmarks in Section 7.2. Although using a standard compiler for all systems is a good solution, it has a drawback. The tool chain is for a machine that does not exist, and it is therefore not readily available and not maintained. This could affect usability in future machines. Seven of the research papers that we surveyed used this benchmark [Mazieres et al. 1999; Muthitacharoen et al. 2002; Padioleau and Ridoux 2003; Santry et al. 1999; Nightingale and Flinn 2004; Cipar et al. 2007; Sobti et al. 2002].

One of the papers [Sobti et al. 2002] further modified the benchmark by removing the *Make* phase and increasing the number of files and directories. Although this removes the complications associated with a compile benchmark and takes care of scalability, data can still be cached depending on the package size. Another paper [Nightingale and Flinn 2004] used Apache for the source files, and measured the time to extract the files from the archive, configure and compile the package, and remove the files. These two papers reported that they used a "modified Andrew benchmark," but since the term "modified" is rather ambiguous, we could not determine if they had used the MAB compiler, or if it was called "modified" because it used a different package or had different phases.

The Andrew benchmark basically combines a compile benchmark and a micro-benchmark.

We suggest using separate compile benchmarks and micro-benchmarks as deemed appropriate (see Section 7.2 and 9 for extensive discussions on each, respectively).

Notable Quotables. We believe some quotations from those papers that used the Andrew benchmark can provide some insight into the reasons for running this benchmark, and for the type of workload that it performs. Six of the fifteen papers that used the Andrew benchmark (or some variant) stated that it was because it was popular or standard. One states, "Primarily because it is customary to do so, we also ran a version of the Andrew benchmark" [Adya et al. 2002]. Six others gave no explicit reason for running the benchmark. The remaining three papers stated that the workload was representative of a user or software developer workload.

Running a benchmark because is it popular or a standard can help readers compare results across papers. Unfortunately, this benchmark has several deficiencies. One paper states that "such Andrew benchmark results do not reflect a realistic workload." [Adya et al. 2002]. Another paper comments that because of the lack of I/O performed, the benchmark "will tend to understate the difference between alternatives." [Kim et al. 2002]. One paper describes that they "modified the benchmark because the 1990 benchmark does not generate much I/O activity by todays standards." [Sobti et al. 2002]. Finally, one paper describes the use of the Andrew benchmark, and how most read requests are satisfied from the cache. "The Andrew Benchmark has been criticized for being old benchmark, with results that are not meaningful to modern systems. It is argued that the workload being tested is not realistic for most users. Furthermore, original Andrew Benchmark used a source tree which is too small to produce meaningful results on modern systems [citation removed]. However, as we stated above, the Benchmarks emphasis on small file performance is still relevant to modern systems. We modified the Andrew Benchmark to use a Linux 2.6.14 source tree [...]. Unfortunately, even with this larger source tree, most of the data by the benchmark can be kept in the OSs page cache. The only phase where file system performance has a significant impact is the copy phase." [Cipar et al. 2007].

It seems that researchers seem to be aware of the benchmark's drawbacks, but still use it because it has become a "standard," because that is what they are accustomed to, or because it is something that other researchers are accustomed to. It is unfortunate that an inadequate benchmark has achieved this status, and we hope that a better option will soon take its place. We believe that FileBench (see Section 10) is promising.

7.4 TPC

The Transaction Processing Performance Council (TPC) is "a non-profit corporation founded to define transaction processing and database benchmarks and to disseminate objective, verifiable TPC performance data to the industry" [Transaction Processing Performance Council 2005]. The organization has strict guidelines about how benchmarks are run, requires full results and configurations to be submitted to them, and audits the results to validate them. To certify benchmark results, companies must have auditors who are accredited by the TPC board standing by throughout the experiments. Whereas this sort of requirement is desirable in a commercial environment, it is not practical for academic papers. Therefore, the benchmarks are used without the strict guidelines attached to them by the TPC. There are four TPC benchmarks currently in use by the database community (TPC-App, TPC-C, TPC-E, and TPC-H). Here we only describe those that were used in the surveyed papers: TPC-B, TPC-C, TPC-D, TPC-H, and TPC-W.

TPC-B. This benchmark has been obsolete since 1995 because it was deemed too simplistic, but was used in one of the surveyed papers in 2005 [Prabhakaran et al. 2005]. Another paper [Denehy et al. 2005] created a benchmark modeled after this workload. The benchmark is designed to stress-test the core of a database system by having several benchmark programs simultaneously submit transactions of a single type as fast as possible. The metric reported is transactions per second.

TPC-C. Created in 1992, this benchmark and adds some complexity that was lacking in older TPC benchmarks, namely TPC-A and TPC-B. It is a data-intensive benchmark portraying the activity of a wholesale supplier where a population of users executes transactions against a database. The supplier has a number of warehouses with stock, and deals with orders and payments. Five different transaction types are used which are either executed immediately or set to be deferred. The database contains nine types of tables with various record and population sizes. The performance metric reported is transactions per minute for TPC-C (tpmC).

TPC-C was used in eight of the surveyed papers [Zhou et al. 2001; Sarkar et al. 2003; Radkov et al. 2004; Huang and Chiueh 2001; Thereska et al. 2004; Ng et al. 2002; Abd-El-Malek et al. 2005; Wachs et al. 2007]. In addition, one paper [Nightingale et al. 2006] used an implementation of TPC-C created by the Open Source Development Lab (OSDL, which was merged into The Linux Foundation in January 2007). The OSDL has developed implementations of several TPC benchmarks [OSDL 2007]. TPC-C is being replaced by TPC-E, which is designed to be representative of current workloads and hardware, is less expensive to run because of more practical storage requirements, and have results that are less dependent on hardware and software configurations.

TPC-D. This benchmark was the precursor to TPC-H (explained next), and has been obsolete since 1999. This is because TPC-D was benchmarking both ad-hoc queries as well as business support and reporting, and could not do both adequately at the same time. TPC-D was split into TPC-H (ad-hoc queries) and TPC-R (business support and reporting).

TPC-H. The workload for this benchmark consists of executing ad-hoc queries against a database and performing concurrent data modifications. Rather than being only data-intensive like TPC-C, this benchmark exercises a larger portion of a database system. It uses queries and data that are relevant to the database community. The benchmark examines large volumes of data, executes queries with a high degree of complexity, and uses the data to give answers to critical business questions (dealing with issues such as supply and demand, profit and revenue, and customer satisfaction). The performance metric reported is called the TPC-H Composite Query-per-Hour Performance Metric (QphH@Size), and reflects multiple aspects of the capability of the system to process queries including the database size, query processing power and throughput. This benchmark was used in three of the surveyed papers [Wachs et al. 2007; Radkov et al. 2004; Gniady et al. 2004].

TPC-W. This benchmark was meant to recreate the workload seen in an Internet commerce environment. This benchmark provides little insight because it is overly complex, difficult to analyze, and is does not recreate the behavior of specific applications [Schmidt et al. 2001]. TPC-W was used in one of the surveyed papers [Huang et al. 2005], but was declared obsolete by TPC approximately 6 months before (in April 2005).

Using a benchmark that is highly regarded and monitored by a council of professionals from the database community certainly adds to its credibility. The benchmarks are kept up-to-date with new version releases, and when serious problems are found with a benchmark, it is declared obsolete. However, none of the papers that used the benchmark had their results audited, and most, if not all, did not run the benchmark according to the specifications. The benchmarks are kept up-to-date with new version releases, and when serious problems are found with a benchmark, it is declared obsolete. A drawback of using TPC benchmarks for performance analysis is that they utilize a database system, which introduces extra complexity. This makes results less comparable between papers and makes the benchmark more difficult to set up. Several papers opted to use traces of the workload instead (see Section 8), and one paper [Lumb et al. 2002] used a synthetic benchmark whose workload was shown to be similar to disk traces of a Microsoft SQL server running TPC-C. Additionally, while most papers specified the database that was used for the experiment, barely any tuning parameters were specified, and none specified the database table layout, which can have dramatic effects on TPC benchmark performance. Although databases are known to have many tuning parameters, one paper specified only two, while others specified only one or none. Some may have used the default settings since they may be less familiar with database systems than file or storage systems, but one paper [Sarkar et al. 2003] specified that the "database settings were fine-tuned for performance" without indicating the exact settings.

7.5 SPEC

The Standard Performance Evaluation Corporation (SPEC), was founded in 1988 by a small number of workstation vendors with the aim of creating realistic, standardized performance tests. SPEC has grown to become a successful performance standardization body with more than 60 member companies [SPEC 2005b].

SFS. The SPEC SFS benchmark [SPEC 2001; Robinson 1999] measures the performance of NFSv2 and v3 servers. It is the official benchmark for measuring NFS server throughput and response time. One of the surveyed papers [Spadavecchia and Zadok 2002] used a precursor to this benchmark, created in 1989, called NFSSTONE [Shein et al. 1989] (not created by SPEC). NFSSTONE performs a series of 45,522 file system operations, mostly executing system calls, to measure how many operations per second an NFS server can sustain. The benchmark performs a mix of operations intended to show typical NFS access patterns [Sandberg et al. 1985]: 53% LOOKUPs, 32% READs, 7.5% READLINKs (symlink traversal), 3.2% WRITEs, 2.3% GETATTRs, and 1.4% CREATEs. This benchmark performs these operations as fast as it can and then reports the average number of operations performed per second, or NFSSTONES. The problems with this benchmark are that it only uses one client so the server is not always saturated, it relied on the client's NFS implementation, and the file sizes and block sizes were not realistic.

Another benchmark, called nhfsstone was developed in 1989 by Legato Systems, Inc. It was similar to NFSSTONE except that instead of the clients executing system calls to communicate with the server, they used packets created by the user-space program. This reduced the dependency on the client, but did not eliminate it, because the client's behavior still depended on the kernel of the machine it was running on.

LADDIS [Watson and Nelson 1992; Wittle and Keith 1993], was created in 1992 by a group of engineers from various companies, and was further developed when SPEC took the project over (this was called SFS 1.0). LADDIS solved some of the deficiencies in the earlier benchmarks by implementing the NFS protocol in user-space, improving the operation mix, allowing multiple clients to generate load on the server simultaneously, providing a consistent method for running the benchmark, and porting the benchmark to several systems. Like the previous benchmarks, given a requested load (measured in ops/sec), LADDIS generates an increasing load of operations and measures the response time until the server is saturated. This is the maximum sustained load that the server can handle under this requested load. As the requested load increases, response time diminishes. The peak throughput is reported.

LADDIS used an outdated workload, and only supported NFSv2 over UDP (no support for NFSv3 or TCP). SFS 2.0 fixed these shortcomings, but several algorithms dealing with request-rate regulation, I/O access, and the file set were found to be defective. SPEC SFS 3.0 fixed the latter, and updated some important features such as the time measurement.

The SFS 3.0 benchmark is run by starting the script on all clients (one will be the main client and direct the others). The number of load-generating processes, the requested load for each run, and the amount of read-ahead and write-behind are specified. For each requested load, SFS reports the average response time. The report is a graph with at least 10 requested loads on the x-axis, and their corresponding response times on the y-axis.

SPEC SFS was used by two of the surveyed papers, one of which ran compliant with SPEC standards [Eisler et al. 2007], and one in which it was not clear [Anderson et al. 2000]. In addition, one paper used a variant of SFS [Patterson et al. 2002], but did not specify how it varied. One issue with SFS is that the number of systems that can be tested is limited to those that speak the NFSv2 and NFSv3 protocols. SFS cannot test changes to NFS clients, and cannot be used to compare an NFS system with a system that speaks another protocol. Whereas this benchmark is very useful for companies that sell filers, its use is limited in the research community. This, combined with the fact that the benchmark is not free (it currently costs $900, or $450 for non-profit or educational institutions), has probably impeded its widespread use in the surveyed papers.

It is of interest to note that the operation mix for the benchmark is fixed. This is good because it standardizes the results more, but also bad because the operation mix can become outdated and may not be appropriate for all settings. Some have claimed that SFS does not resemble any NFS workload they have observed, and that each NFS trace that they examined had unique characteristics, raising the question of whether one can construct a standard workload [Zhu et al. 2005]. One can change the operations mix for SFS, so in some sense it can be used as a workload generator. However, the default, standard operations mix must be used to report any standard results that can be compared with other systems. It seems that this may be the end for the SFS benchmark, because NFSv4 is already being deployed, and SPEC has not stated any plans to release a new version of SFS. In addition, Spencer Shepler, one of the chairs of the NFSv4 IETF Working Group, has stated that SPEC SFS is "unlikely to be extended to support NFSv4," and that FileBench (see Section 10) will probably be used instead [Shepler 2005].

SDM. The SPEC SDM benchmarking suite [SPEC 2004] was made in 1991 and produces a workload that simulates a software development environment with a large number of users. It contains two sub-benchmarks, 057.SDET and 061.Kenbus1, both of which feed randomly ordered scripts to the shell with commands like `make`, `cp`, `diff`, `grep`, `man`, `mkdir`, `spell`, etc. Both use a large number of concurrent processes to generate significant file system activity.

The measured metric is the number of scripts completed per hour. One script is generated for each "user" before the timing begins, and each contains separate subtasks executed in a random order. Each user is then given a home directory that is populated with the appropriate directory tree and files. A shell is started for each user which uses its own execution script. The timer is stopped when all scripts have completed execution.

There are two main differences between these two benchmarks. First, Kenbus simulates users typing characters at a rate of three characters per second as opposed to SDET which reads as fast as possible. Second, the command set used in SDET is a lot richer and many of the commands do a lot more work than in Kenbus.

The SDET benchmark was used in two surveyed papers [Ng et al. 2002; Seltzer et al. 2000] to measure file system performance, but not all commands executed by the benchmark exercise the file system. It is meant to measure the performance of the system as a whole, and not any particular subsystem. In addition, the benchmark description states that it exercises the `tmp` directories heavily, which means that either the benchmark needs to be changed, or the file system being tested must be mounted as the system disk. However, the benchmark does give a reasonable idea of how a file system would affect everyday workloads. This benchmark is currently being updated and is being called SMT (System MultiTasking) [SPEC 2003], with the main goal of ensuring that all systems perform the same amount of work regardless of configuration. However, the SMT description has not been updated since 2003, so it is unknown as to whether or not it will be deployed.

Viewperf. The SPECviewperf benchmark [SPEC 2007] is designed to measure the performance of a graphics subsystem, and was used in one of the surveyed research papers [Gniady et al. 2004]. However, we will not delve into this benchmark's details because it is inappropriate to use it as a file system benchmark, as it exercises many parts of the OS other than the file system.

Web99. The SPECweb99 benchmark [SPEC 2005a], replaced by SPECweb2005 in 2005, is used for evaluating the performance of Web servers. One of the surveyed research papers [Nightingale et al. 2006] used it. Its workload is comprised of dynamic and static GET operations, as well as POST operations. We omit further discussion because it is a Web server benchmark, and was used by the surveyed paper to specifically measure their file system using a network-intensive workload.

7.6 SPC

The Storage Performance Council (SPC) [SPC 2007] develops benchmarks focusing on storage subsystems. Its goal is to have more vendors use industry-standard benchmarks, and publish those results in a standard way. The council consists of several major vendors in the storage industry, as well as some academic institutions. The SPC currently has two benchmarks available to its members: SPC-1 and SPC-2. Neither were used in the surveyed papers, but are clearly noteworthy.

SPC-1. This benchmark's workload is designed to perform typical functions of business-critical applications. The workload is comprised of predominately random I/O operations, and performs both queries and update operations. This type of workload is typical of on-line transaction processing (OLTP) systems, database systems, or mail server applications. SPC-1 is designed to accurately measure performance and price/performance on both direct attach or network storage subsystems. It includes three tests.

The first test phase, called "Primary Metrics," has three phases. In the first, throughput sustainability is tested for three hours at steady state. The second phase lasts for ten minutes, and tests the maximum attainable throughput in I/Os per second. The third phase lasts fifty minutes, and maps the relationship between response time and throughput by measuring latencies at various load levels, defined as percentages of the throughput achieved in the previous phase. The third phase also determines the optimal average response time of a lightly-loaded storage configuration.

The second test was designed to prove that the maximum I/O request throughput results that were determined in the first test are repeatable and reproducible. It does this by running similar but shorter workloads as the first test to collect the same metrics. In the third and final test, SPC-1 demonstrates that the system provides non-volatile and persistent data storage. It does this by writing random data to random locations over the total capacity of the storage system for at least ten minutes. The writes are recorded in a log. The system is shut down, and caches that employ battery backup are flushed or emptied. The system is the restarted, and the written data is verified.

SPC-2. This benchmark is characterized by a predominantly sequential workload (in contrast to SPC-1's random workload). The workload is intended to demonstrate the performance of business-critical applications that require large-scale, sequential data transfers. Such applications include large file processing (scientific computing, large-scale financial processing), large database queries (data mining, business intelligence), and on-demand video. SPC-2 includes four tests and it measures the throughput.

The first test checks data persistence, similar to the third test of SPC-1. The second test measures large file processing. It has three phases (write-only, read-write, read-only), each consisting of two run sequences, each composed of five runs (thirty runs in total). Each run consists of a certain sized transfer with a certain number of streams. The third test is the "large database query test," which has two phases (1,024 KiB transfer size and 64 KiB transfer size). Each phase consists of two run sequences (four outstanding requests and one outstanding request), and each sequence consists of five runs where the number of streams is varied (ten runs total). The fourth and final test is the "video on demand delivery test," in which several streams of data are transferred.

Since the Storage Performance Council has many prominent storage vendors as members, it is likely that its benchmarks will be widely used in industry. However, their popularity in academia is yet to be seen, as the benchmarks are currently only available to SPC members, or for $500 for non-member academic institutions. Of course, academic institutions will probably not follow all of the strict benchmark guidelines and pay the expensive result filing fees, but the benchmarks would still allow for good comparisons.

7.7 NetNews

The NetNews benchmark [Swartz 1996], created in 1996, is a shell script that performs a small-file workload comparable to that which is seen on a USENET NetNews server. It performs some setup work, and then executes the following three phases multiple times:

(1) *Unbatch* - Measures the receiving and storing of new articles. Enough data is used to ensure that all caches are flushed.

(2) *Batch* - Measures the sending of backlogged articles to other sites. Articles are batched on every third pass to reduce the runtime of the benchmark, but they should be

large enough so they will not be cached when they are re-used.

(3) *Expire* - Measures the removal of “expired” articles. Since article timestamps are not relevant in a benchmark setting, the number of history entries is recorded after each *unbatch* phase. This information is used to expire articles from more than some number of previous passes. The list of articles to be deleted is written to one file while the modified history file is written to another.

One of the surveyed papers [Seltzer et al. 2000] used this benchmark without the Batch phase to analyze performance. It is a good benchmark in the sense that it is meta-data intensive, and stresses the file system (given a large enough workload). However, while the data size used in the surveyed paper was considered by the authors to be large (270MB for the *unbatch* phase, and 250MB for the *expire* phase), it is much smaller than sizes seen in the real world. The paper states that “two years ago, a full news feed could exceed 2.5 GB of data, or 750,000 articles per day. Anecdotal evidence suggests that a full news feed today is 15–20 GB per day.” This shows how a static workload from 1996 is not realistic just four years later. In addition to the unrealistic workload size, a USENET NetNews server workload is not very common these days, and it may be difficult to extrapolate results from this benchmark to applications that more people use.

7.8 Other Macro-Benchmarks

This section describes two infrequently used macro-benchmarks that appeared in the surveyed research papers.

NetBench and dbench. NetBench [VeriTest 2002] is a benchmark used to measure the performance of file servers. It uses a network of PCs to generate file I/O requests to a file server. According to the dbench README file [A. Tridgell 1999], NetBench’s main drawback is that properly running the benchmark requires a lab with 60–150 PCs running Windows, connected with switched fast Ethernet, and a high-end server. It also states that since the benchmark is “very fussy,” the machines should be personally monitored. Because of these factors, this benchmark is rarely used outside of the corporate world.

Dbench is an open-source program that runs on Linux machines and produces the same file system load as NetBench would on a Samba server but without making any networking calls. It was used in one surveyed paper to measure performance [Schmuck and Haskin 2002]. The metrics reported by dbench are true throughput and the throughput expected on a Win9X machine.

Because there is no source code or documentation for NetBench, we were limited to analyzing the dbench source code. The dbench program is run with one parameter—the number of clients to simulate. It begins by creating a child process for each client. There is a file that uses commands from a Windows trace of NetBench, which each child process reads one line at a time. The benchmark executes the Linux equivalent of each Windows command from the trace, and the file is processed repeatedly for ten minutes. The main process is signaled every second to calculate and print the throughput up to that point, and to signal the child processes to stop when the benchmark is over. The first two minutes of the run is a warm-up period, and statistics collected during this time are kept separate from statistics collected during the remainder of the run.

The main question with dbench is how closely it approximates NetBench. There are three problems that we have discovered. First, there is no one-to-one correspondence between Windows and Linux operations, and the file systems natively found on each of these

OSs are quite different, so it is unknown how accurate the translation is. Second, NetBench has each client running on a separate machine, while Linux has the main process and all the child processes running on the same machine. The added computation needed for all of the process management plus the activities of each process may affect the results. In addition, if there are many concurrent processes, the benchmark may be analyzing the performance of other subsystems, such as the scheduler, more than the file system. Third, dbench processes the trace file repeatedly and does not consider timings. The source of the trace file is not clear, and so it is unknown how well the operations in the trace reflect the NetBench workload. Caching of the trace file may also affect the results since it is processed multiple times by several clients on the same machine.

8. REPLAYING TRACES

Traces are logs of operations that are collected, and later replayed to generate the same workload (if done correctly). They have the same goal as macro-benchmarks have (see Section 7): to produce a workload which represents a real-world environment. However, although it may be uncertain whether or not a macro-benchmark succeeds at this, a trace will definitely recreate the workload that was traced if it is captured and played back correctly. One must ensure, however, that the captured workload is representative of the system's intended real-world environment. Nineteen of the surveyed papers used traces as part of their performance analysis [Adya et al. 2002; Sivathanu et al. 2004; Nightingale and Flinn 2004; Weil et al. 2006; Rowstron and Druschel 2001; Arpaci-Dusseau et al. 2003; Zhang et al. 2002; Lu et al. 2002; Lumb et al. 2003; Dimitrijevic et al. 2003; Flinn et al. 2003; Sivathanu et al. 2004; Schindler et al. 2004; Tolia et al. 2004; Weddle et al. 2007; Peterson et al. 2007; Tian et al. 2007; Prabhakaran et al. 2005; Tolia et al. 2003].

Some use traces of machines running macro-benchmarks such as TPC-C [Dimitrijevic et al. 2003; Zhang et al. 2002], TPC-H [Schindler et al. 2004], or compile benchmarks [Weil et al. 2006]. These traces differ from the norm in that they are traces of a synthetic workload rather than a real-world environment. It is unclear why a trace of a compile benchmark was used, rather than running the compile benchmark itself. However, since the actual TPC benchmarks require a database system and have rather complicated setups, papers may opt to use traces of the TPC benchmark instead. However, it is important to replay the trace in a similar environment as where the trace was gathered. For example, one paper [Dimitrijevic et al. 2003] used a trace of a TPC run that used a file-system–based database, but replayed it in an environment that bypassed the file system to access the block device directly. For more information on TPC benchmarks, see Section 7.4.

There are four problem areas with traces today, described next: the capture method, the replay method, trace realism, and trace availability.

Capture method. There is no accepted way to capture traces, and this can be a source of confusion. Traces can be captured at the system call, VFS, networking, and driver levels.

The most popular way is to capture traces at the system-call level primarily because it is easy and the system call API is portable [W. Akkerman 2002; Mummert and Satyanarayanan 1994; Ousterhout et al. 1985]. One benefit of capturing at the system-call level is that it does not differentiate between requests that are satisfied from the cache and those that are not. This allows one to test changes in caching policies. An important drawback of capturing at the system call level is that memory-mapped operations cannot be captured. Traces captured at the VFS level contain cached and non-cached requests, as well as memory-mapped requests. However, VFS tracer portability is limited even between different versions of the same OS. Existing VFS tracers are available for Linux [Aranya et al. 2004] and Windows NT [Vogels 1999; Roselli et al. 2000].

Network-level traces contain only requests that were not satisfied from the cache. Network-level capturing is only suitable for network file systems. Network packet traces can be collected using specialized devices or software tools like *tcpdump*. Specialized tools can capture and pre-process only the network file systems related packets [Ellard and Seltzer 2003a; Blaze 1992]. Driver-level traces contain only non-cached requests and cannot correlate the requests with the associated meta-data without being provided or inferring additional information. For example, read requests to file meta-data and data read requests cannot be easily distinguished [Ruemmler and Wilkes 1993].

The process by which the trace is captured must be explained, and should be distributed along with the trace if others will be using it. In five of the surveyed papers, the authors captured their own traces, but four of them did not specify how this was done.

When collecting file system traces for studies, many use tracing tools that are customized for a single study. These systems are built either in an ad-hoc manner [Ousterhout et al. 1985; Roselli et al. 2000], modify standard tools [Roselli et al. 2000; Ellard et al. 2003], or are not well documented in research texts. Their emphasis is on studying the characteristics of file system operations and not on developing a systematic or reusable infrastructure for tracing. Often, the traces excluded useful information for others conducting new studies; information excluded could concern the initial state of the machines or hardware on which the traces were collected, some file system operations and their arguments, pathnames, and more.

Replay method. Replaying a file system trace correctly is not as easy as it may appear. Before one can start replaying, the trace itself may need to be modified. Any missing operations have to be guessed so that operations that originally succeeded do not fail (and those that failed should not succeed) [Zhu et al. 2005]. For example, files must be created before they are accessed. In addition, the trace may be scaled spatially or temporally [Zhu et al. 2005]. For parallel applications, finding inter-node dependencies and inter-I/O compute times in the trace improves replay correctness [Mesnier et al. 2007]. Once the trace is ready, the target file system must be prepared. Files that are assumed to exist in the trace must exist on the file system with at least the size of the largest offset accessed in it. This will ensure that the trace can be replayed, but the resulting file system will have no fragmentation, and will only include files that were accessed in the trace, which is not realistic. The solution here is to age the file system. Of course, the aging method should be described, because the replay will now differ from the original run, as well as from replays that aged the file system differently. There are fewer issues with preparing block-level traces for replay on traditional storage devices, since block accesses generally do not have context associated with them. In this case, missing operations cannot be known, and aging the storage system will not affect the behavior of the benchmark, since the blocks specified by the trace will be accessed regardless of the storage system's previous state.

It is natural to replay traces at the level the traces were captured. However, replaying file system traces at the system call level makes it impossible to replay high I/O-rate traces with the same speed as they were captured on the same hardware. This is because replaying adds overheads associated with system calls (context switches, verifying arguments, copies between user-space and kernel-space [Anderson et al. 2004]). VFS-level replaying requires kernel-mode development, but can use the time normally spent on executing system calls to prefetch and schedule future events [Joukov et al. 2005]. Network-level replaying is popular because it can be done entirely from the user-level. Unfortunately, it is only applicable to network file systems. Driver-level replaying allows one to control the physical data position on the disk, and is often done from user-level.

Replay speed is an important consideration, and is subject to some debate. Some believe that the trace should be replayed with the original timings, although none of the surveyed papers specified that they did this. There are replaying tools, such as Buttress [Anderson et al. 2004], that have been shown to follow timings accurately. However, with the advent of faster machines, it would be unreasonable to replay an older trace with the same timings. On the other hand, if the source of the trace was a faster machine, it may not be possible to use the same timings.

There is another school of thought that believes that the trace should be played as fast as possible, ignoring the timings. Five of the surveyed papers did so [Flinn et al. 2003; Tolia et al. 2004; Peterson et al. 2007; Nightingale and Flinn 2004; Prabhakaran et al. 2005]. Any trace replay speed will measure something slightly different than what the original system's behavior when the trace was being captured. However, replaying a trace as fast as possible changes the behavior more than other speeds do due to factors such as caching, read-ahead, and interactions with page write-back and other asynchronous events in the OS. It assumes an I/O bottleneck, and ignores operation dependencies.

A compromise between using the original timings and ignoring the timings is to play back the trace with some speedup factor. Three of the surveyed papers [Zhang et al. 2002; Sivathanu et al. 2004; Weddle et al. 2007] replayed with both the original timings as well as at increased speeds. By doing so they were able to observe the effects of increasing the pressure on the system. Although this is better than replaying at only one speed, it is not clear what scaling factors to choose.

We believe that currently the best option is to replay the trace as fast as possible and report the average operations per second. However, it is crucial to respect dependencies in file system traces, and not simply run one operation after the other. For example, TBBT [Zhu et al. 2005] has one replay mode called "conservative order," which sends out a request only after all previous operations have completed, and another called "FS dependency order," which applies some file system orderings to the operations (for example, it will not write to a file before it has received a reply that the file was created).

An ideal way of recreating a traced workload would be to first normalize the think times using both the hardware and software (e.g., OS, libraries, relevant programs) specifications of the system that produced the trace, and then to calibrate it using specifications of the machine being used to replay the trace. How to do this accurately is still an open question, and the best we can do right now is take the results with a grain of salt.

However replaying is done, the method should be clearly stated along with the results. Of the 19 surveyed papers that utilized tracing for performance analysis, we found that 15 did not specify the tool used to replay and 11 did not specify the speed. Not specifying the replay tool hinders the reader from judging the accuracy of the replay. For those that did not specify the speed, one can guess that traces were replayed as fast as possible since this is common and easy, but it is still possible that other timings were used.

Realism. Whether or not a trace accurately portrays the intended real-world environment is important to consider. One aspect of this problem is traces becoming stale. For traces whose age we could determine, the average age was 5.7 years, with some as old as 11 years. Studies have shown that characteristics such as file sizes, access patterns, file system size, file types, and directory size distribution have changed over the years [Roselli et al. 2000; Agrawal et al. 2007].

Additionally, the trace should ideally be of many users executing similar workloads (preferably unaware that the trace is being collected so that their behavior is not biased). For example, one of the surveyed papers [Adya et al. 2002] used a one-hour–long trace of one developer as the backbone of their evaluation. Such a small sample may not accurately represent a whole population.

Trace availability. Researchers may collect traces themselves, request them directly from other researchers, or obtain them from some third party that makes traces publicly available. Reusing traces, where appropriate, can encourage comparisons between papers and allow results to be reproduced. However, traces can become unavailable for several reasons, some of which are discussed here. One surveyed paper [Rowstron and Druschel 2001] used traces that are available via FTP, but the company that captures and hosts the traces states that they remove traces after seven days. In some cases, those who captured the traces have moved on and are unavailable. Additionally, since trace files are usually very large (traces from HP Labs, which are commonly used, can be up to 9GB), researchers may not save them for future use. Some traces are even larger (approximately 1TB, compressed 20–30:1), so even if the authors still have the trace, they may insist on transferring it by shipping a physical disk. To resolve these types of issues, traces should be stored in centralized, authoritative repositories of traces for all to use. In 2003, SNIA created a technical working group called IOTTA (I/O Traces, Tools & Analysis) to attack this problem. They have established a world-wide trace repository, with several traces in compatible formats and with all of the necessary tools [SNIA 2007]. There are also two smaller repositories hosted by universities [LASS 2006; PEL 2001].

Privacy and anonymization is a concern when collecting and distributing traces, but it should be done while not harming the usability of the traces. For example, one could encrypt sensitive fields, each with a different encryption key. Different mappings for each field remove the possibility of correlation between related fields. For example, UID = 0 and GID = 0 usually occur together in traces, but this cannot be easily inferred from the anonymized traces in which the two fields have been encrypted using different keys. Keys could be given out in private to decrypt certain fields if desired [Aranya et al. 2004]. Although this does not hide all important information, such as the number of users on the system, it should provide enough privacy for most scenarios.

9. MICRO-BENCHMARKS

In this section we describe the micro-benchmarks that were used in the surveyed research papers, and reflect on their positive qualities, drawbacks, and how appropriate they were in the context of the papers that they appeared in.

In contrast to the macro-benchmarks described in Section 7, micro-benchmark workloads usually consist of a small number of types of operations and serve to highlight some specific aspect of the file system.

We discuss Bonnie and Bonnie++ in Section 9.1, the Sprite benchmarks in Section 9.2, ad-hoc micro-benchmarks in Section 9.3, and using system utilities to create workloads in Section 9.4.

9.1 Bonnie and Bonnie++

Bonnie, developed in 1990, performs a series of tests on a single file, which is 100KB by default [Bray 1996]. For each test, Bonnie reports the number of bytes processed per elapsed second, the number of bytes processed per CPU second, and the percent of CPU usage (user and system). The tests are:

Sequential output. The file is created one character at a time, and then recreated one 8KB chunk at a time. Each chunk is then read, dirtied, and rewritten.

Sequential input. The file is read once one character at a time, and then again one chunk at a time.

Random seeks. A number of processes seek to random locations in the file and read a chunk of data. The chunk is modified and rewritten 10% of the time. The documentation states that the default number of processes is four, but we have checked the source code and only three are created (this shows the benefit of having open-source benchmarks). The default number of seeks for each process is 4,000 and is one of several hard-coded values.

Even though this is a fairly well-known benchmark [Bryant et al. 2001], of all the papers that we surveyed only one of them [Zadok et al. 2001] used it. Care must be taken to ensure that the working file size is larger than the amount of memory on the system so that not all read requests are satisfied from the page cache. The Bonnie documentation recommends using a file size that is at least 4 times bigger than the amount of available memory. However, the biggest file size that was used in the surveyed research paper was equal to the amount of memory on the machine. The small number of papers using Bonnie may be due to the three drawbacks it has.

First, unlike Postmark (see Section 7.1), Bonnie does not use a single pseudo-random number generator for all OSs. This injects some variance between benchmarks run on different OSs, and results may not be comparable.

Second, the options are not parameterized [Bryant et al. 2001]. Of all of the values mentioned above, only the file size is configurable from the command line. The rest of the values are hard-coded in the program. In addition, Bonnie does not allow the workload to be fully customized. A mix of sequential and random access is not possible, and the number of writes can never exceed the number of reads (because a write is done only after a chunk is read and modified).

Third, reading and writing one character at a time tests the library call throughput more than the file system because the function that Bonnie calls (`getc`) uses buffering.

Bonnie++ [Coker 2001] was created in 2000 and used in one of the surveyed papers [Peterson et al. 2005]. It differs from Bonnie in three ways. First, it is written in C++ rather than C. Second, it uses multiple files to allow accessing data sets that are larger than 2GB. Third, it adds several new tests that benchmark the performance of `create`, `stat`, and `unlink`. Although it adds some useful features to Bonnie, Bonnie++ still suffers from the same three drawbacks of Bonnie.

9.2 Sprite LFS

Two micro-benchmarks from the 1992 Sprite LFS file system [Rosenblum 1992] are sometimes used in research papers for performance analysis: the large file benchmark and the small file benchmark.

Sprite LFS Large File Benchmark. Three papers [Mazieres et al. 1999; Wang et al. 1999; Wang et al. 2002] included this benchmark in their performance evaluation. The benchmark has five phases:

(1) Create a 100MB file using sequential writes.
(2) Reads the file sequentially.
(3) Writes 100MB randomly to the existing file.
(4) Reads 100MB randomly from the file.
(5) Reads the file sequentially.

The most apparent fault with this benchmark is that it uses a fixed-size file and therefore does not scale. It also uses the `random` library routine rather than having a built-in pseudo-random number generator (as Postmark does—see Section 7.1) which may make results incomparable across machines with different implementations.

Another fault is that the caches are not cleaned between each phase, and so some number of operations in a given phase may be serviced from the cache (the amount would actually depend on the pseudo-random number generator for phases 3 and 4, further emphasizing the need for a common generator). However, two of the papers did specify that caches were cleaned after each write phase. It should be noted that for the random read phase, the benchmark ends up reading the entire file, and so the latency of this phase depends on the file system's read-ahead algorithm (and the pseudo-random number generator).

One of the good points about the benchmark is that each stage is timed separately, with a high level of accuracy, and only relevant portions of the code are timed. For example, when performing random writes, it first generates the random order (though it uses a poorly-designed algorithm which is $O(N^2)$ in the worst case), and then starts timing the writes.

Sprite LFS Small File Benchmark. This benchmark was used in six papers [Wang et al. 2002; Denehy et al. 2005; Mazieres et al. 1999; Kaminsky et al. 2003; Li et al. 2004; Wang et al. 1999]. It has three phases:

(1) Creating 10,000 1KB files by creating and opening a file, writing 1KB of data, and closing the file.
(2) Reading the files.
(3) Deleting the files.

Some papers varied the number of files and their sizes, and two specified that the caches were flushed after the write phase. We could not obtain the source code for this benchmark, so it seems that each author may rewrite it because it is so simple. This would make it difficult to compare results across papers because the source code may be different. To show this, we have developed five versions of the code:

LFS-SH1. A `bash` script that creates the files by coping data from `/dev/zero` using the `dd` program (with a block size of 1 byte, and count of 1,024), reads the files using `cat`, and deletes them with `rm`.

LFS-SH2. Similar to LFS-SH1, but `dd` uses a 1,024-byte block size, and a count of 1.

LFS-SH3. The same as LFS-SH1, but uses `cp` instead of `dd` to create the files.

LFS-C. A C implementation.

LFS-PL. A Perl implementation.

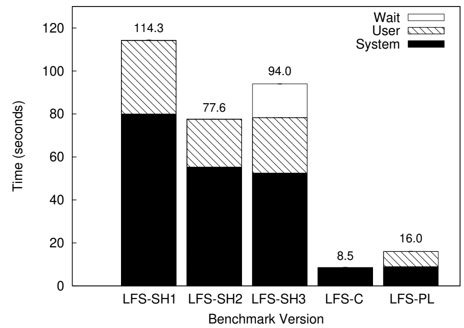

Fig. 6. Time taken to execute various versions of the Sprite LFS small file benchmark. Note that error bars are shown, but are small and difficult to see.

The source code for all five versions is available at `www.fsl.cs.sunysb.edu/project-fsbench.html`. The results, shown in Figure 6, clearly demonstrate that different implementations yield significantly different results. The three `bash` script versions are much slower than the others because every operation in the benchmark forks a new process. In addition, all of the implementations except for LFS-SH3 have insignificant wait time components, showing that file system activity is minimal.

9.3 Ad-Hoc Micro-Benchmarks

Until now, we have been discussing widely-available benchmarks in isolation. In contrast, ad-hoc benchmarks are written by the authors for in-house use. In this section, we describe ad-hoc micro-benchmarks in the context of the papers that they appear in, since the benchmarks alone are usually not very interesting. 62 of the 106 surveyed research papers used ad-hoc micro-benchmarks for at least one of their experiments (191 total ad-hoc micro-benchmarks).

These benchmarks all have a general drawback. Because they are not made available to other researchers, they are not reproducible. Even if they are described in detail (which is usually not the case), another implementation will certainly differ (see Section 9.2 for experimental evidence). In addition, since these benchmarks are not widely used, they are not tested as much as widely available benchmarks, and therefore are more prone to bugs. One good aspect about these benchmarks is that we have noticed that usually some reasoning behind the benchmark is described.

Because micro-benchmarks test the performance of a file system under very specific circumstances, they should not be used alone to describe general performance characteristics. The only exception to this rule is if a minor change is made in some code and there is a clear explanation as to why no other execution paths are affected.

We have identified three reasonable ways of using micro-benchmarks. The first acceptable way of using ad-hoc micro-benchmarks is to better understand the results of other benchmarks. For example, one paper [Ng et al. 2002] measured the throughput of reads and writes for sequential, random, and identical block access for this purpose. Because the results are not meant to be compared across papers, the reproducibility is no longer much of an issue. Sequential access may show the best case because of the short disk head seeks and predictable nature, which is where read-ahead and prefetching shine. The random read micro-benchmark is generally used to measure read throughput in a scenario where there is no observable access pattern and disk head seeks are common. This type of behavior has been observed in database workloads. The random aspect of the benchmark inherently inhibits its reproducibility across machines, as discussed in Section 7.1. Since this paper involved network storage, repeatedly reading from the same block gives the time for issuing a request over the network that results in a cache hit. The reason for writing to the same block, however, was not explained. Nevertheless, using ad-hoc micro-benchmarks to explain other results is a good technique. Another paper [Fu et al. 2000] used the read phases of the LFS benchmarks to examine performance bottlenecks.

A second method of using these benchmarks is to use several ad-hoc micro-benchmarks to analyze the performance of a variety of operations, as eight papers did. This can provide a sense of how the system would perform compared to some baseline for commonly used operations, and may allow readers to estimate the overheads for other workloads.

The third way is to use the micro-benchmark to isolate a specific aspect of the system. For example, a tracing file system used an ad-hoc micro-benchmark to exercise the file system by producing large traces that general-purpose benchmarks such as Postmark could not produce, thereby showing worst-case performance [Aranya et al. 2004]. Another used a simple sequential read benchmark to illustrate different RPC behavior [Magoutis et al. 2003]. Others used ad-hoc micro-benchmarks to show how the system behaves under those specific conditions. Most of these papers focused on the read, write, and stat operations, varying the access patterns, the number of threads, and the number of files. However, most did not use the micro-benchmarks to show worst-case behavior.

In addition, ad-hoc micro-benchmarks can be used in the initial phases of benchmarking to explore the behavior of a system. This can provide useful data about code that requires optimization or to make decisions about what additional benchmarks would most effectively show the system's behavior.

9.4 System Utilities

Some papers use standard utilities to create workloads instead of creating workloads from scratch, as discussed in Section 9.3. Some examples of benchmarks of this type that were used in the surveyed papers are:

—`wc` [Meter and Gao 2000], `grep` [Meter and Gao 2000; Fraser and Chang 2003]: sequentially reading from one or more files.

—`cp` [Zadok et al. 2001; Padioleau and Ridoux 2003; Schindler et al. 2002; Santry et al. 1999; Muniswamy-Reddy et al. 2004]: sequentially reading from one file while copying to another.

—`diff` [Schindler et al. 2002]: sequentially reading two files and comparing them.

—`tar` [Lee et al. 1999; Anderson et al. 2000], `gzip` [DeBergalis et al. 2003]: sequentially reading from a file and appending to a set of files while consuming CPU.

Using these utilities is slightly better than creating ad-hoc benchmarks because these utilities are widely-available and there is no misunderstanding about what the benchmark does. However, none of the papers specified what version they were using, which could lead to some (possibly minor) changes in workloads. For example, different versions of `grep` use different I/O strategies. However, researchers can easily specify tool version and eliminate all ambiguity. However, an important flaw in using these utilities is that the benchmarks do not scale, and depend on the input files which are not standardized.

10. CONFIGURABLE WORKLOAD GENERATORS

Configurable workload generators generally have lower flexibility when compared to creating custom benchmarks, but they require less setup time and are usually more reproducible. In addition, since they are more widely-used and established than ad-hoc benchmarks, it is likely that they contain fewer bugs. We discuss some of the more popular generators here.

Iometer. This workload generator and measurement tool was developed by Intel in 1998 and originally developed for Windows [OSDL 2004]. Iometer was given to the Open Source Development Lab in 2001, which open-sourced and ported it to other OSs. The authors claim that it can be configured to emulate the disk or network I/O load of any program or benchmark, and that it can be used to generate entirely synthetic I/O loads. It can generate and measure loads on single or multiple (networked) systems. Iometer can be used for measurement and characterization of disk and network controller performance, bus latency and bandwidth, network throughput to attached drives, shared bus performance, and hard drive and network performance.

The parameters for configuring tests include the following: the run time; the amount of time to run the benchmark before collecting statistics (useful for making sure the system is in a "steady state"); the number of threads; the number of targets (i.e., disks or network interfaces); the number of outstanding I/O operations; the workload to run.

The parameters for a thread's workload include: the percent of transfers that are a given size; the ratio of reads to writes; the ratio of random to sequential accesses; number of transfers in a burst; time to wait between bursts; the alignment of each I/O on the disk; the size of the reply, if any, to each I/O request. The test also includes a large selection of metrics to use when displaying results, and can save and load configuration files.

Iometer has four qualities not found in many other benchmarks. First, it scales well, since the user inputs the amount of time the test should run, rather than the amount of work to be performed. Second, allowing the system to reach steady state is a good practice, although it may be more useful to find this point by statistical methods rather than by trusting the user to input a correct time. Third, it allows for configuration files to be easily distributed and publicized by saving the configuration file so that benchmarks can be run with exactly the same workloads. Although researchers can publicize parameters for other benchmarks, there is no standard format so some parameters are bound to be left unreported. Fourth, having a suite that runs multiple tests with varying parameters saves time and reduces errors. However, there are tools such as Auto-pilot [Wright et al. 2005] that can automate benchmarks with greater control (for example, the machine can reboot automatically between runs, run helper scripts, etc.).

A drawback of Iometer is that it does not leave enough room for customization. Although it can recreate most commonly-used workloads, hard-coding the possibilities for workload specification and performance metrics reduces its flexibility. For example, the percentage of reads that are random is not sufficient to describe all read patterns. One read pattern suggested for testing read-ahead is reading blocks from a file in the following patterns: 1, 51, 101,...; 1, 2, 51, 52, 101, 102,...; 1, 2, 3, 51, 52, 53, 101, 102, 103,... [Tang 1995]. Such patterns cannot be recreated with Iometer.

Surprisingly, even though Iometer has many useful features, only two papers used it [Sarkar et al. 2003; Yu et al. 2000]. This may be because Iometer was unable to generate the desired workload, as described above. However, most workloads are fairly straightforward, so this is less of a factor. More likely, researchers simply may not know about it, or be familiar with it. This may be why researchers prefer to write their own micro-benchmarks rather than using a workload generator. Furthermore, it does not seem that the ability to save configuration files improved the reporting of workloads in the research papers that used Iometer (neither paper fully described their micro-benchmarks).

Buttress. The goal of Buttress is to issue I/O requests with a high accuracy, even when high throughputs are requested [Anderson et al. 2004]. This is important for Buttress' trace replay capability, as well for obtaining accurate inter-I/O times for its workload generation capability. Accurate I/O issuing is not present in most benchmarks, and the authors show how important it is to have. The Buttress toolkit issues read and write requests close to their intended issue time, can achieve close to maximum possible throughput, and can replay I/O traces as well as generate synthetic I/O patterns. Several interesting techniques were employed to ensure these properties. In addition, the toolkit is flexible (unlike Iometer) because users can specify their own workloads using a simple event-based programming interface. However, this also makes it more difficult to reproduce benchmarks from other papers (it is easier to specify simple parameters, as with Iometer). Two of the surveyed research papers, both from HP labs, have used this toolkit [Lumb et al. 2003; Lu et al. 2002]. Unfortunately, Buttress is only available by special request from HP.

FileBench. This workload generator from Sun Microsystems [R. McDougall and J. Mauro 2005] is configured using a scripting language. FileBench includes scripts that can generate application-emulating or micro-benchmark workloads, and users may write their own scripts for custom benchmarks. The application workloads it currently emulates are an NFS mail server (similar to Postmark—see Section 7.1), a file server (similar to SPEC SFS—see Section 7.5), a database server, a Web server, and a Web proxy.

FileBench also generates several micro-benchmark workloads, some of which are similar to Bonnie (see Section 9.1) or the copy phase of the Andrew benchmark (see Section 7.3). In addition to the workloads that come with FileBench, it has several useful features: (1) workload scripts can easily be reused and published, (2) the ability to choose between multiple threads or multiple processes, (3) micro-second accurate latency and cycle counts per system call, (4) thread synchronization, (5) warm-up and cool-down phases to measure steady-state activity, (6) configurable directory structures, (7) database emulation features (e.g., semaphores, synchronous I/O, etc.).

Only one of the surveyed papers [Gulati et al. 2007] used FileBench, possibly because Filebench was made publicly available only in 2005 on Solaris, and a Linux port was created soon after. However, it is highly configurable and it is possible that researchers will be able to use it for running many of their benchmarks.

Fstress. This workload generator has similar parameters to other generators [Andrerson 2002]. One can specify the following: the distributions for file, directory, and symlink counts; the maximum directory tree depth; popularity in accesses for newly created objects; file sizes; operation mix; I/O sizes; and load level. Like SPEC SFS (see Section 7.5), it only runs over NFSv3, and constructs packets directly rather than relying on a client implementation. However, NFSv3 is currently being replaced by NFSv4, so supporting this protocol would be necessary to ensure relevance. Like Iometer, there are limited workload configuration parameters. Another drawback is that requests are sent at a steady rate, so bursty I/O patterns cannot be simulated. Fstress was not used in any of the surveyed papers.

11. BENCHMARKING AUTOMATION

Proper benchmarking is an iterative process for many reasons. In our experience, there are four primary reasons for this. First, when running a benchmark against a given configuration, you must run each test a sufficient number of times to gain confidence that your results are accurate. Second, most software does not exist in a vacuum—there is at least one other related system or a system that serves as a baseline for comparison. In addition to your own system, you must benchmark the other systems and compare your performance to those. Third, benchmarks often expose bugs or inefficiencies in your code, which require changes. After fixing these bugs (or simply adding new features), you must re-run your benchmarks. Fourth, after doing a fair number of benchmarks, you inevitably run into unexpected, anomalous, or just interesting results. To explain these results, you often need to change configuration parameters or measure additional quantities—necessitating additional iterations of your benchmark. Therefore, it is natural to automate the benchmarking process from start to finish.

Auto-pilot [Wright et al. 2005] is a suite of tools that we developed for producing accurate and informative benchmark results. We have used Auto-pilot for over five years, on dozens of projects. As each project is slightly different, we continuously enhanced Auto-pilot and increased its flexibility for each one. The result is a stable and mature package that saves days and weeks of repetitive labor on each project. Auto-pilot consists of four major components: a tool to execute a set of benchmarks described by a simple configuration language, a collection of sample shell scripts for file system benchmarking, a data extraction and analysis tool, and a graphing tool. The analysis tool can perform all of the statistical tests that we described in Section 3.

12. EXPERIMENTAL EVALUATIONS

In this section we describe the methods we used to benchmark file systems to show some of their qualities. Our goal is to demonstrate some of the common pitfalls of file system benchmarks. We describe the file system that we used for the benchmarks in Section 12.1. In Section 12.2, we show some of the faults that exist in commonly used benchmarks.

12.1 Slowfs

To reveal some characteristics of various benchmarks, we have modified the ext2 file system to slow down certain operations. We call this file system *Slowfs*. Rather than calling the normal function for an operation, we call a new function which does the following:

(1) $start$:= getcc() [get current time in CPU cycles]

(2) Calls the original function for the operation

(3) now := getcc()

(4) $goal$:= now + ((now - $start$) * 2^N) - (now - $start$)

(5) while (getcc() $\leq$ goal) { schedule() }

The net effect of this is a slow-down of a factor of 2^N for the operation. The operations to slow down and N are given as mount-time parameters. For this article we slowed down the following operations:

—*read.* Reads data from the disk.

—*prepare write and commit write.* Used by the file system to write data to disk for the write system call. We refer to these operations collectively as WRITE for the remainder of the article.

—*lookup.* Takes a directory and a file name, and returns an in-memory inode.

If no operation was slowed down, we call it EXT2. If all of the above operations were slowed down we call it ALL. We experimented with the above three functions because they are among the most common found in benchmarks. Note that this type of slow-down exercises the CPU and not I/O, and that a slow-down of a certain factor is as seen inside the file system, not by the user (the amount of overhead as seen by the user varies with each benchmark). For example, heavy use of the CPU can be found in file systems that perform encryption, compression, or checksumming for integrity or duplicate elimination.

The source code for Slowfs is available at `www.fsl.cs.sunysb.edu/project-fsbench.html`.

12.2 Hiding Overheads

In this section we use Slowfs to prove some of the claims that we have made in this article.

Compile benchmarks. In our first experiment, we compared Slowfs and ext2 for configuring and compiling OpenSSH versions 3.5, 3.7, and 3.9. We used Slowfs with the read operation slowed down by several factors. The results are shown in Figure 7. Only the results for versions 3.5 and 3.7 are shown, because they showed the highest overheads. We chose to slow down the read operation because, as shown in Section 7.2, it is the most time-consuming operation for this benchmark. Because a compile benchmark is CPU-intensive, such extraordinary overheads as a factor of 32 on read can go unnoticed (the factor of 32 comes from setting N to 5, as described in Section 12.1). For all of these graphs, the half-widths were less than 1.5% of the mean, and the $CPU\%$ was always more than 99.2%, where $CPU\% = \frac{time_{user}+time_{system}}{time_{elapsed}} \times 100$. In the following discussion, we do not include user or I/O times because they were always either statistically indistinguishable or very close (these two values were not affected by the Slowfs modifications).

For the configure phase, the highest overhead was 2.7% for elapsed time, and 10.1% for system time (both for version 3.7). For the compile phase, the highest overhead was 4.5% for elapsed time and 59.2% for system time (both for version 3.5). Although 59.2% is a noticeable overhead, this can be hidden by only reporting the elapsed time overhead.

We also conducted the same compile benchmarks with slowing down each of the operations listed in Section 12.1 only by a factor of five. We slowed them down separately as well as together. There was *no* statistical difference between ext2 and any of these slowed down configurations.

These results clearly show that even with extraordinary delays in critical file system operations, compile benchmarks show only marginal overheads because they are bound by CPU time spent in user-space. As mentioned in Section 7.3, the deficiencies in compile benchmarks apply to the Andrew benchmark as well.

Postmark. We tested Slowfs with three different Postmark configurations (described in Table IV). The FSL configuration is the one we have been using in our laboratory [Aranya et al. 2004; Wright et al. 2003], the CVFS configuration is from the CVFS research pa-

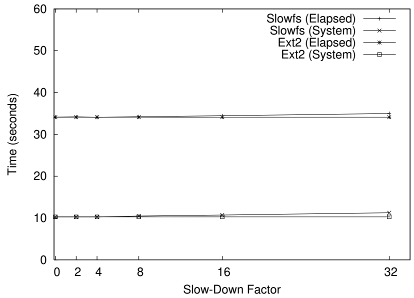

(a) Configure phase for version 3.5

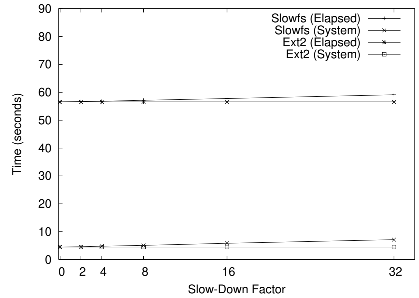

(b) Compile phase for version 3.5

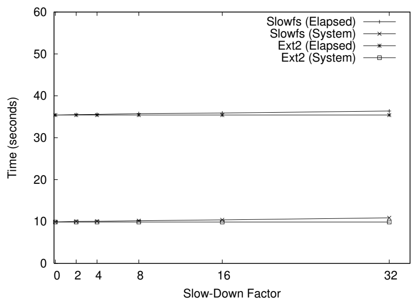

(c) Configure phase for version 3.7

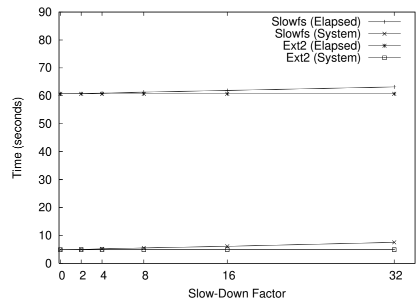

(d) Compile phase for version 3.7

Fig. 7. Time taken to configure and compile OpenSSH versions 3.5 and 3.7 on ext2 and on Slowfs with the read operation slowed down by several factors. Note the different scales for the Y-axes. The half-widths were always less than 1.5% of the mean.

Table IV. Postmark configurations used in our Slowfs experiments.

Parameter	**FSL**	**CVFS**	**CVFS-LARGE**
Number of Files	20,000	5,000	5,000
Number of Subdirectories	200	50	50
File Sizes	512 bytes–10KB	512 bytes–10KB	512–328,072 bytes
Number of Transactions	200,000	20,000	20,000
Operation Ratios	equal	equal	equal
Read Size	4KB	4KB	4KB
Write Size	4KB	4KB	4KB
Buffered I/O	no	no	no

per [Soules et al. 2003], and CVFS-LARGE is similar to the CVFS configuration, but we used the median size of a mailbox on our campus's large mail server for the file size. We used a similar configuration before [Muniswamy-Reddy et al. 2004], but have updated the file size. We used Postmark version 1.5, and used Slowfs to slow down each of the operations separately, as well as together, by a factor of four. The results are shown in Figure 8.

The graphs show us two important features of this benchmark. First, if we look at the EXT2 bar in each graph, we can see how much changing the configurations can effect the results. The three are very different, and are clearly incomparable (FSL takes over 55 times longer than CVFS, and CVFS-LARGE is still almost twice as long as FSL). Second, we can

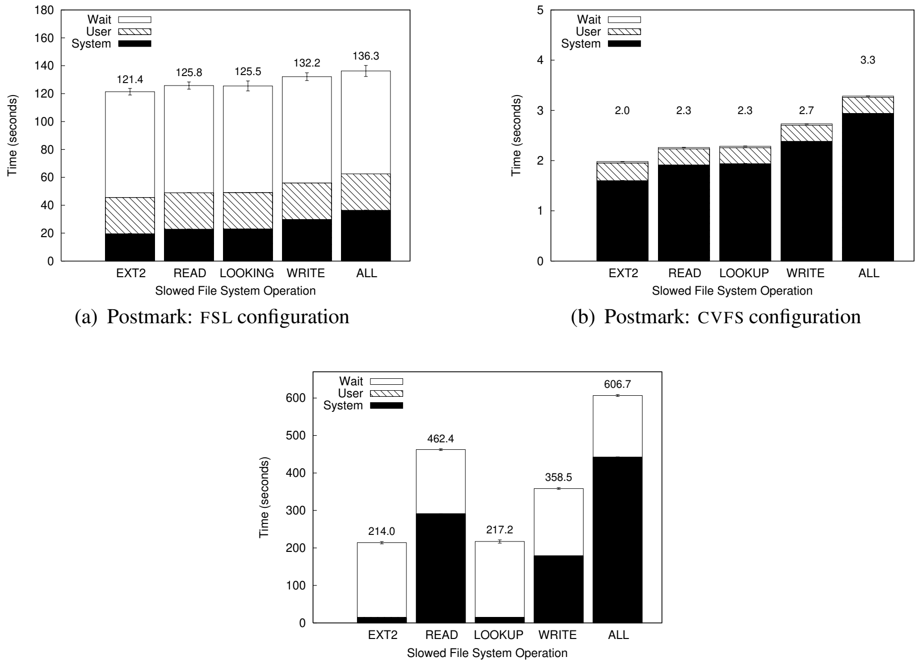

(a) Postmark: FSL configuration

(b) Postmark: CVFS configuration

(c) Postmark: CVFS-LARGE configuration

Fig. 8. Time taken to execute the Postmark benchmark with several configurations while slowing down various file system operations using Slowfs. Note the different scales for the Y-axes.

see that different configurations show the effects of Slowfs in varying degrees.

For example, slowing down reads yields an elapsed time overhead of 3.6% for FSL (16.7% system time), 14.1% for CVFS (19.4% system time), and 116% for CVFS-LARGE (2,055% system time) over ext2 (EXT2). We can see that in the CVFS configuration, there is no wait time on the graph. This is because the configuration was so small that the benchmark finished before the flushing daemon could write to the disk. CVFS has larger overheads than FSL because writes are a smaller component of the benchmark, and so reads become a larger component. CVFS-LARGE has higher overheads than the other two configurations because it has much larger files, and so there is more data to be read. Similarly, when all operations are slowed down (ALL), there is an elapsed time overhead of 12.3% for FSL (85.6% system time), 65.8% for CVFS (83.5% system time), and 183% for CVFS-LARGE (3,177% system time).

Depending on the characteristics of the file system being tested, it is possible to choose a configuration that will yield low overheads. Even so, we see that Postmark sufficiently exercises the file system and shows us meaningful overheads as long as the workload is large enough to produce I/O (i.e., the working set is larger than available memory and the benchmark runs for enough time). This is in contrast to the compile benchmarks, which barely show any overheads.

13. CONCLUSIONS

We have examined a range of file system and storage benchmarks and described their positive and negative qualities, with the hope of furthering the understanding of how to choose appropriate benchmarks for performance evaluations. We have done this by surveying 106 file-system and storage-related research papers from a selection of recent conferences and by conducting our own experiments. We also advised on how benchmarks should be run and how results should be presented. This advice was summarized in our suggested guidelines (see Section 3).

We suggest that with the current set of available benchmarks, the most accurate method of conveying a file or storage system's performance is by using at least one macro-benchmark or a trace, as well as several micro-benchmarks. Macro-benchmarks and traces are intended to give an overall idea of how the system would perform under some workload. If traces are used, then special care must be taken with regard to how they are captured, how they are replayed, and how closely they resemble the intended real-world workload. In addition, micro-benchmarks should be used to help understand the system's performance, test multiple operations to provide a sense of overall performance, or highlight interesting features about the system (such as cases where it performs particularly well or poor).

Performance evaluations must improve in their descriptions of *what* they did, as well as *why* they did it, which is equally important. Explaining the reasoning behind one's actions is an important principle in research, but seems to be ignored in some file system and storage performance evaluations. Ideally, there should be some analysis of the system's expected behavior, and various benchmarks either proving or disproving the hypotheses. This provides more insight into the behavior than just a graph or table can.

We believe that the current state of performance evaluations as seen in the surveyed research papers is bleak. Computer science is still a relatively young field, and the experimental evaluations needs to move further in the direction of precise science. One part of the solution is that standards clearly need to be raised. This will have to be done both by reviewers putting more emphasis on a system's evaluation, and by researchers by raising the bar. Another part of the solution is that researchers need to be better informed. We hope that this paper, and our continuing work, will help researchers understand the problems that exist with file and storage system benchmarking. The final aspect of the solution to this problem is creating standardized benchmarks, or benchmarking suites, based on open discussion among file system and storage researchers.

We believe that future research can help alleviate the situation by answering questions such as:

(1) How can we accurately portray various real-world workloads?

(2) How can we accurately compare results from benchmarks that were run on different machines and systems and at different times?

To help answer the first question, we need a method of determining how close two workloads are to each other. To answer the second, we believe that benchmark results can be normalized for the machine they were run on. In order to standardize benchmarks, we feel that there is a need to have a group such as the SPC to standardize and maintain file system benchmarks, and for the SPC benchmarks to be more widely used by the storage community. We are currently working on some of these problems and there is still much work to be done, but we hope that with time the situation will improve.

The project Web site (`www.fsl.cs.sunysb.edu/project-fsbench.html`) contains the data collected for this survey, our suggestions for proper benchmarking techniques, and the source code and machine configurations we used in the experiments throughout the paper.

ACKNOWLEDGMENTS

We would like to thank all the people who helped review earlier drafts of this work, as well as attendees of the FAST 2005 BoF for their valuable comments. We would also like to thank the members of the File Systems and Storage Laboratory at Stony Brook for their helpful comments and advice.

This work was made partially possible thanks to NSF awards CNS-0133589 (CAREER), CCR-0310493 (CyberTrust), CNS-0614784 (CSR), and CCF-0621463 (HECURA)—as well as two HP/Intel gifts numbers 87128 and 88415.1.

REFERENCES

A. Tridgell. 1999. dbench-3.03 README. `http://samba.org/ftp/tridge/dbench/README`.

Abd-El-Malek, M., II, W. V. C., Cranor, C., Ganger, G., Hendricks, J., Klosterman, A. J., Mesnier, M., Prasad, M., Salmon, B., Sambasivan, R. R., Sinnamohideen, S., Strunk, J. D., Thereska, E., Wachs, M., and Wylie, J. J. 2005. Ursa Minor: Versatile Cluster-based Storage. In *Proceedings of the 4th USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 59–72.

Adya, A., Bolosky, W. J., Castro, M., Cermak, G., Chaiken, R., Douceur, J. R., Howell, J., Lorch, J. R., Theimer, M., and Wattenhofer, R. P. 2002. FARSITE: Federated, Available, and Reliable Storage for an Incompletely Trusted Environment. In *Proceedings of the 5th Symposium on Operating System Design and Implementation*. USENIX Association, Boston, MA, 1–14.

Agrawal, N., Bolosky, W. J., Douceur, J. R., and Lorch, J. R. 2007. A five-year study of file-system metadata. In *Proceedings of the 5th USENIX Conference on File and Storage Technologies*. USENIX Association, San Jose, CA, 31–45.

Aguilera, M. K., Ji, M., Lillibridge, M., MacCormick, J., Oertli, E., Andersen, D., Burrows, M., Mann, T., and Thekkath, C. A. 2003. Block-Level Security for Network-Attached Disks. In *Proceedings of the 2nd USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 159–174.

Anderson, D. C., Chase, J. S., and Vahdat, A. M. 2000. Interposed Request Routing for Scalable Network Storage. In *Proceedings of the 4th Usenix Symposium on Operating System Design and Implementation*. USENIX Association, San Diego, CA, 259–272.

Anderson, E., Hobbs, M., Keeton, K., Spence, S., Uysal, M., and Veitch, A. 2002. Hippodrome: Running Circles Around Storage Administration. In *Proceedings of the 1st USENIX Conference on File and Storage Technologies*. USENIX Association, Monterey, CA, 175–188.

Anderson, E., Kallahalla, M., Uysal, M., and Swaminathan, R. 2004. Buttress: A Toolkit for Flexible and High Fidelity I/O Benchmarking. In *Proceedings of the 3rd USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 45–58.

Andrerson, D. 2002. Fstress: A flexible network file service benchmark. Tech. Rep. TR-2001-2002, Duke University. May.

Aranya, A., Wright, C. P., and Zadok, E. 2004. Tracefs: A File System to Trace Them All. In *Proceedings of the 3rd USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 129–143.

Arpaci-Dusseau, A. C., Arpaci-Dusseau, R. H., Burnett, N. C., Denehy, T. E., Engle, T. J., Gunawi, H. S., Nugent, J. A., and Popovici, F. I. 2003. Transforming Policies into Mechanisms with Infokernel. In *Proceedings of the 19th ACM Symposium on Operating Systems Principles*. ACM SIGOPS, Bolton Landing, NY, 90–105.

Blaze, M. 1992. NFS Tracing by Passive Network Monitoring. In *Proceedings of the USENIX Winter Conference*. USENIX Association, San Francisco, CA.

BRAY, T. 1996. The Bonnie home page. www.textuality.com/bonnie.

BRYANT, R., FORESTER, R., AND HAWKES, J. 2002. Filesystem Performance and Scalability in Linux 2.4.17. In *Proceedings of the Annual USENIX Technical Conference, FREENIX Track*. USENIX Association, Monterey, CA, 259–274.

BRYANT, R., RADDATZ, D., AND SUNSHINE, R. 2001. PenguinoMeter: A New File-I/O Benchmark for Linux. In *Proceedings of the 5th Annual Linux Showcase & Conference*. USENIX Association, Oakland, CA, 5–10.

CHEN, P. M. AND PATTERSON, D. A. 1993. A New Approach to I/O Performance Evaluation - Self-Scaling I/O Benchmarks, Predicted I/O Performance. In *Proceedings of the 1993 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*. ACM SIGOPS, Seattle, WA, 1–12.

CIPAR, J., CORNER, M. D., AND BERGER, E. D. 2007. Tfs: A transparent file system for contributory storage. In *Proceedings of the 5th USENIX Conference on File and Storage Technologies*. USENIX Association, San Jose, CA, 215–229.

COKER, R. 2001. The Bonnie++ home page. www.coker.com.au/bonnie++.

CORBETT, P., ENGLISH, B., GOEL, A., GRCANAC, T., KLEIMAN, S., LEONG, J., AND SANKAR, S. 2004. Row-Diagonal Parity for Double Disk Failure Correction. In *Proceedings of the 3rd USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 1–14.

DABEK, F., KAASHOEK, M. F., KARGER, D., AND MORRIS, R. 2001. Wide-Area Cooperative Storage with CFS. In *Proceedings of the 18th ACM Symposium on Operating Systems Principles*. ACM, Banff, Canada.

DEBERGALIS, M., CORBETT, P., KLEIMAN, S., LENT, A., NOVECK, D., TALPEY, T., AND WITTLE, M. 2003. The Direct Access File System. In *Proceedings of the 2nd USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 175–188.

DENEHY, T. E., ARPACI-DUSSEAU, A. C., AND ARPACI-DUSSEAU, R. H. 2002. Bridging the information gap in storage protocol stacks. In *Proceedings of the Annual USENIX Technical Conference*. USENIX Association, Monterey, CA, 177–190.

DENEHY, T. E., ARPACI-DUSSEAU, A. C., AND ARPACI-DUSSEAU, R. H. 2005. Journal-guided Resynchronization for Software RAID. In *Proceedings of the 4th USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 87–100.

DIMITRIJEVIC, Z., RANGASWAMI, R., AND CHANG, E. 2003. Design and Implementation of Semi-preemptible IO. In *Proceedings of the 2nd USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 145–158.

EISLER, M., CORBETT, P., KAZAR, M., NYDICK, D. S., AND WAGNER, J. C. 2007. Data ontap gx: A scalable storage cluster. In *Proceedings of the 5th USENIX Conference on File and Storage Technologies*. USENIX Association, San Jose, CA, 139–152.

ELLARD, D., LEDLIE, J., MALKANI, P., AND SELTZER, M. 2003. Passive NFS Tracing of Email and Research Workloads. In *Proceedings of the 2nd USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA.

ELLARD, D. AND SELTZER, M. 2003a. New NFS Tracing Tools and Techniques for System Analysis. In *Proceedings of the Annual USENIX Conference on Large Installation Systems Administration*. USENIX Association, San Diego, CA.

ELLARD, D. AND SELTZER, M. 2003b. NFS Tricks and Benchmarking Traps. In *Proceedings of the Annual USENIX Technical Conference, FREENIX Track*. USENIX Association, San Antonio, TX, 101–114.

FLINN, J., SINNAMOHIDEEN, S., TOLIA, N., AND SATYANARYANAN, M. 2003. Data Staging on Untrusted Surrogates. In *Proceedings of the 2nd USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 15–28.

FRASER, K. AND CHANG, F. 2003. Operating System I/O Speculation: How Two Invocations Are Faster Than One. In *Proceedings of the Annual USENIX Technical Conference*. USENIX Association, San Antonio, TX, 325–338.

FU, K., KAASHOEK, M. F., AND MAZIÈRES, D. 2000. Fast and Secure Distributed Read-Only File System. In *Proceedings of the 4th Usenix Symposium on Operating System Design and Implementation*. USENIX Association, San Diego, CA, 181–196.

GHEMAWAT, S., GOBIOFF, H., AND LEUNG, S. T. 2003. The Google File System. In *Proceedings of the 19th ACM Symposium on Operating Systems Principles*. ACM SIGOPS, Bolton Landing, NY, 29–43.

GNIADY, C., BUTT, A. R., AND HU, Y. C. 2004. Program-Counter-Based Pattern Classification in Buffer Caching. In *Proceedings of the 6th Symposium on Operating Systems Design and Implementation*. ACM SIGOPS, San Francisco, CA, 395–408.

GOPAL, B. AND MANBER, U. 1999. Integrating Content-based Access Mechanisms with Hierarchical File Systems. In *Proceedings of the 3rd Symposium on Operating Systems Design and Implementation*. ACM SIGOPS, New Orleans, LA, 265–278.

GRÖNVALL, B., WESTERLUND, A., AND PINK, S. 1999. The Design of a Multicast-based Distributed File System. In *Proceedings of the 3rd Symposium on Operating Systems Design and Implementation*. ACM SIGOPS, New Orleans, LA, 251–264.

GULATI, A., NAIK, M., AND TEWARI, R. 2007. Nache: Design and implementation of a caching proxy for nfsv4. In *Proceedings of the 5th USENIX Conference on File and Storage Technologies*. USENIX Association, San Jose, CA, 199–214.

HOWARD, J. H., KAZAR, M. L., MENEES, S. G., NICHOLS, D. A., SATYANARAYANAN, M., SIDEBOTHAM, R. N., AND WEST, M. J. 1988. Scale and performance in a distributed file system. *ACM Transactions on Computer Systems 6,* 1 (February), 51–81.

HUANG, H., HUNG, W., AND SHIN, K. 2005. FS2: Dynamic Data Replication in Free Disk Space for Improving Disk Performance and Energy Consumption. In *Proceedings of the 20th ACM Symposium on Operating Systems Principles*. Brighton, UK, 263–276.

HUANG, L. AND CHIUEH, T. 2001. Charm: An i/o-driven execution strategy for high-performance transaction processing. In *Proceedings of the Annual USENIX Technical Conference*. USENIX Association, Boston, MA, 275–288.

JOGLEKAR, A., KOUNAVIS, M. E., AND BERRY, F. L. 2005. A scalable and high performance software iscsi implementation. In *Proceedings of the 4th USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 267–280.

JOUKOV, N., TRAEGER, A., IYER, R., WRIGHT, C. P., AND ZADOK, E. 2006. Operating System Profiling via Latency Analysis. In *Proceedings of the 7th Symposium on Operating Systems Design and Implementation*. ACM SIGOPS, Seattle, WA, 89–102.

JOUKOV, N., WONG, T., AND ZADOK, E. 2005. Accurate and efficient replaying of file system traces. In *Proceedings of the 4th USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 337–350.

KALLAHALLA, M., RIEDEL, E., SWAMINATHAN, R., WANG, Q., AND FU, K. 2003. Plutus: Scalable Secure File Sharing on Untrusted Storage. In *Proceedings of the 2nd USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 29–42.

KAMINSKY, M., SAVVIDES, G., MAZIERES, D., AND KAASHOEK, M. F. 2003. Decentralized User Authentication in a Global File System. In *Proceedings of the 19th ACM Symposium on Operating Systems Principles*. ACM SIGOPS, Bolton Landing, NY.

KATCHER, J. 1997. PostMark: A New Filesystem Benchmark. Tech. Rep. TR3022, Network Appliance. `www.netapp.com/tech_library/3022.html`.

KIM, J. M., CHOI, J., KIM, J., NOH, S. H., MIN, S. L., CHO, Y., AND KIM, C. S. 2000. A Low-Overhead, High-Performance Unified Buffer Management Scheme That Exploits Sequential and Looping References. In *Proceedings of the 4th Usenix Symposium on Operating System Design and Implementation*. USENIX Association, San Diego, CA, 119–134.

KIM, M., COX, L., AND NOBLE, B. 2002. Safety, Visibility, and Performance in a Wide-Area File System. In *Proceedings of the 1st USENIX Conference on File and Storage Technologies*. USENIX Association, Monterey, CA.

KROEGER, T. M. AND LONG, D. D. E. 2001. Design and implementation of a predictive file prefetching algorithm. In *Proceedings of the Annual USENIX Technical Conference*. USENIX Association, Boston, MA, 105–118.

LASS. 2006. UMass Trace Repository. `http://traces.cs.umass.edu`.

LEE, Y., LEUNG, K., AND SATYANARAYANAN, M. 1999. Operation-based update propagation in a mobile file system. In *Proceedings of the Annual USENIX Technical Conference*. USENIX Association, Monterey, CA, 43–56.

LI, J., KROHN, M., MAZIÈRES, D., AND SHASHA, D. 2004. Secure Untrusted Data Repository (SUNDR). In *Proceedings of the 6th Symposium on Operating Systems Design and Implementation*. ACM SIGOPS, San Francisco, CA, 121–136.

LU, C., ALVAREZ, G. A., AND WILKES, J. 2002. Aqueduct: Online Data Migration with Performance Guarantees. In *Proceedings of the 1st USENIX Conference on File and Storage Technologies*. USENIX Association, Monterey, CA.

LUMB, C. R., MERCHANT, A., AND ALVAREZ, G. A. 2003. Façade: Virtual Storage Devices with Performance Guarantees. In *Proceedings of the 2nd USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 131–144.

LUMB, C. R., SCHINDLER, J., AND GANGER, G. R. 2002. Freeblock Scheduling Outside of Disk Firmware. In *Proceedings of the 1st USENIX Conference on File and Storage Technologies*. USENIX Association, Monterey, CA, 275–288.

MACCORMICK, J., MURPHY, N., NAJORK, M., THEKKATH, C., AND ZHOU, L. 2004. Boxwood: Abstractions as the Foundation for Storage Infrastructure. In *Proceedings of the 6th Symposium on Operating Systems Design and Implementation*. ACM SIGOPS, San Francisco, CA, 105–120.

MAGOUTIS, K., ADDETIA, S., FEDOROVA, A., AND SELTZER, M. I. 2003. Making the Most Out of Direct-Access Network Attached Storage. In *Proceedings of the 2nd USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 189–202.

MAGOUTIS, K., ADDETIA, S., FEDOROVA, A., SELTZER, M. I., CHASE, J. S., GALLATIN, A. J., KISLEY, R., WICKREMESINGHE, R. G., AND GABBER, E. 2002. Structure and Performance of the Direct Access File System. In *Proceedings of the Annual USENIX Technical Conference*. USENIX Association, Monterey, CA.

MAZIÉRES, D. 2001. A Toolkit for User-Level File Systems. In *Proceedings of the Annual USENIX Technical Conference*. USENIX Association, Boston, MA, 261–274.

MAZIERES, D., KAMINSKY, M., KAASHOEK, M. F., AND WITCHEL, E. 1999. Separating key management from file system security. In *Proceedings of the 17th ACM Symposium on Operating Systems Principles*. ACM, Charleston, SC, 140–153.

MEMIK, G., KANDEMIR, M., AND CHOUDHARY, A. 2002. Exploiting Inter-File Access Patterns Using Multi-Collective I/O. In *Proceedings of the 1st USENIX Conference on File and Storage Technologies*. USENIX Association, Monterey, CA.

MESNIER, M. P., WACHS, M., SAMBASIVAN, R. R., LOPEZ, J., HENDRICKS, J., GANGER, G. R., AND O'HALLARON, D. 2007. //trace: Parallel trace replay with approximate causal events. In *Proceedings of the 5th USENIX Conference on File and Storage Technologies*. USENIX Association, San Jose, CA, 153–167.

METER, R. V. 1997. Observing the Effects of Multi-Zone Disks. In *Proceedings of the Annual USENIX Technical Conference*. USENIX Association, Anaheim, CA, 19–30.

METER, R. V. AND GAO, M. 2000. Latency Management in Storage Systems. In *Proceedings of the 4th Usenix Symposium on Operating System Design and Implementation*. USENIX Association, San Diego, CA, 103–118.

MILLER, E., FREEMAN, W., LONG, D., AND REED, B. 2002. Strong Security for Network-Attached Storage. In *Proceedings of the 1st USENIX Conference on File and Storage Technologies*. USENIX Association, Monterey, CA, 1–13.

MOGUL, J. 1999. Brittle Metrics in Operating Systems Research. In *Proceedings of the IEEE Workshop on Hot Topics in Operating Systems (HOTOS)*. ACM, Rio Rica, AZ, 90–95.

MUMMERT, L. AND SATYANARAYANAN, M. 1994. Long term distributed file reference tracing: Implementation and experience. Tech. Rep. CMU-CS-94-213, Carnegie Mellon University, Pittsburgh, PA.

MUNISWAMY-REDDY, K., HOLLAND, D. A., BRAUN, U., AND SELTZER, M. 2006. Provenance-aware storage systems. In *Proceedings of the Annual USENIX Technical Conference*. USENIX Association, Boston, MA, 43–56.

MUNISWAMY-REDDY, K., WRIGHT, C. P., HIMMER, A., AND ZADOK, E. 2004. A Versatile and User-Oriented Versioning File System. In *Proceedings of the 3rd USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 115–128.

MUTHITACHAROEN, A., CHEN, B., AND MAZIERES, D. 2001. A Low-Bandwidth Network File System. In *Proceedings of the 18th ACM Symposium on Operating Systems Principles*. ACM, Banff, Canada.

MUTHITACHAROEN, A., MORRIS, R., GIL, T. M., AND CHE, B. 2002. Ivy: A Read/Write Peer-to-Peer File System. In *Proceedings of the 5th Symposium on Operating System Design and Implementation*. USENIX Association, Boston, MA, 31–44.

NG, W. T., SUN, H., HILLYER, B., SHRIVER, E., GABBER, E., AND OZDEN, B. 2002. Obtaining High Performance for Storage Outsourcing. In *Proceedings of the 1st USENIX Conference on File and Storage Technologies*. USENIX Association, Monterey, CA, 145–158.

NIGHTINGALE, E. B., CHEN, P., AND FLINN, J. 2005. Speculative Execution in a Distributed File System. In *Proceedings of the 20th ACM Symposium on Operating Systems Principles*. Brighton, UK, 191–205.

NIGHTINGALE, E. B. AND FLINN, J. 2004. Energy-Efficiency and Storage Flexibility in the Blue File System. In *Proceedings of the 6th Symposium on Operating Systems Design and Implementation*. ACM SIGOPS, San Francisco, CA, 363–378.

NIGHTINGALE, E. B., VEERARAGHAVAN, K., CHEN, P. M., AND FLINN, J. 2006. Rethink the Sync. In *Proceedings of the 7th Symposium on Operating Systems Design and Implementation*. ACM SIGOPS, Seattle, WA, 1–14.

NUGENT, J., ARPACI-DUSSEAU, A., AND ARPACI-DUSSEAU, R. 2003. Controlling Your PLACE in the File System with Gray-box Techniques. In *Proceedings of the Annual USENIX Technical Conference*. USENIX Association, San Antonio, TX, 311–323.

OSDL. 2004. Iometer Project. `www.iometer.org/`.

OSDL. 2007. Database Test Suite. `www.osdl.org/lab_activities/kernel_testing/osdl_database_test_suite/`.

OUSTERHOUT, J. 1990. Why aren't operating systems getting faster as fast as hardware? In *Proceedings of the Summer USENIX Technical Conference*. USENIX, Anaheim, CA, 247–256.

OUSTERHOUT, J., COSTA, H., HARRISON, D., KUNZE, J., KUPFER, M., AND THOMPSON, J. 1985. A Trace-Driven Analysis of the UNIX 4.2 BSD File System. In *Proceedings of the 10th ACM Symposium on Operating System Principles*. ACM, Orcas Island, WA, 15–24.

PADIOLEAU, Y. AND RIDOUX, O. 2003. A Logic File System. In *Proceedings of the Annual USENIX Technical Conference*. USENIX Association, San Antonio, TX, 99–112.

PAPATHANASIOU, A. E. AND SCOTT, M. L. 2004. Energy efficient prefetching and caching. In *Proceedings of the Annual USENIX Technical Conference*. USENIX Association, Boston, MA, 255–268.

PATTERSON, H., MANLEY, S., FEDERWISCH, M., HITZ, D., KLEINMAN, S., AND OWARA, S. 2002. SnapMirror: File System Based Asynchronous Mirroring for Disaster Recovery. In *Proceedings of the 1st USENIX Conference on File and Storage Technologies*. USENIX Association, Monterey, CA, 117–129.

PEEK, D. AND FLINN, J. 2006. EnsemBlue: Integrating Distributed Storage and Consumer Electronics. In *Proceedings of the 7th Symposium on Operating Systems Design and Implementation*. ACM SIGOPS, Seattle, WA, 219–232.

PEL. 2001. BYU Trace Distribution Center. `http://tds.cs.byu.edu/tds`.

PETERSON, Z. N. J., BURNS, R., ATENIESE, G., AND BONO, S. 2007. Design and implementation of verifiable audit trails for a versioning file system. In *Proceedings of the 5th USENIX Conference on File and Storage Technologies*. USENIX Association, San Jose, CA, 93–106.

PETERSON, Z. N. J., BURNS, R., J. HERRING, A. S., AND RUBIN, A. D. 2005. Secure Deletion for a Versioning File System. In *Proceedings of the 4th USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 143–154.

PRABHAKARAN, V., AGRAWAL, N., BAIRAVASUNDARAM, L. N., GUNAWI, H. S., ARPACI-DUSSEAU, A. C., AND ARPACI-DUSSEAU, R. H. 2005. IRON File Systems. In *Proceedings of the 20th ACM Symposium on Operating Systems Principles*. Brighton, UK, 206–220.

PRABHAKARAN, V., ARPACI-DUSSEAU, A. C., AND ARPACI-DUSSEA, R. H. 2005. Analysis and Evolution of Journaling File Systems. In *Proceedings of the Annual USENIX Technical Conference*. USENIX Association, Anaheim, CA, 105–120.

QUINLAN, S. AND DORWARD, S. 2002. Venti: a new approach to archival storage. In *Proceedings of the 1st USENIX Conference on File and Storage Technologies*. USENIX Association, Monterey, CA, 89–101.

R. MCDOUGALL AND J. MAURO. 2005. FileBench. `www.solarisinternals.com/si/tools/filebench/`.

RADKOV, P., YIN, L., GOYAL, P., SARKAR, P., AND SHENOY, P. 2004. A Performance Comparison of NFS and iSCSI for IP-Networked Storage. In *Proceedings of the 3rd USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 101–114.

RHEA, S., EATON, P., GEELS, D., WEATHERSPOON, H., ZHAO, B., AND KUBIATOWICZ, J. 2003. Pond: The OceanStore Prototype. In *Proceedings of the 2nd USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 1–14.

ROBINSON, D. 1999. The advancement of NFS benchmarking: SFS 2.0. In *Proceedings of the 13th USENIX Systems Administration Conference*. USENIX Association, Seattle, WA, 175–185.

ROSELLI, D., LORCH, J. R., AND ANDERSON, T. E. 2000. A Comparison of File System Workloads. In *Proc. of the Annual USENIX Technical Conference*. USENIX Association, San Diego, CA, 41–54.

ROSENBLUM, M. 1992. The design and implementation of a log-structured file system. Ph.D. thesis, Electrical Engineering and Computer Sciences, Computer Science Division, University of California.

ROWSTRON, A. AND DRUSCHEL, P. 2001. Storage Management and Caching in PAST, A Large-scale, Persistent Peer-to-peer Storage Utility. In *Proceedings of the 18th ACM Symposium on Operating Systems Principles*. ACM, Banff, Canada.

RUEMMLER, C. AND WILKES, J. 1993. UNIX Disk Access Patterns. In *Proceedings of the Winter USENIX Technical Conference*. USENIX Association, San Diego, CA, 405–420.

RUWART, T. M. 2001. File system performance benchmarks, then, now, and tomorrow. In *Proceedings of the 14th IEEE Symposium on Mass Storage Systems*. IEEE, San Diego, CA.

SAITO, Y., KARAMANOLIS, C., KARLSSON, M., AND MAHALINGAM, M. 2002. Taming Aggressive Replication in the Pangaea Wide-Area File System. In *Proceedings of the 5th Symposium on Operating System Design and Implementation*. USENIX Association, Boston, MA, 15–30.

SANDBERG, R., GOLDBERG, D., KLEIMAN, S., WALSH, D., AND LYON, B. 1985. Design and implementation of the Sun Network Filesystem. In *Proceedings of the Summer USENIX Technical Conference*. USENIX Association, Portland, Oregon, 119–130.

SANTRY, D. S., FEELEY, M. J., HUTCHINSON, N. C., VEITCH, A. C., CARTON, R. W., AND OFIR, J. 1999. Deciding When to Forget in the Elephant File System. In *Proceedings of the 17th ACM Symposium on Operating Systems Principles*. ACM, Charleston, SC, 110–123.

SARKAR, P., UTTAMCHANDANI, S., AND VORUGANTI, K. 2003. Storage Over IP: When Does Hardware Support Help? In *Proceedings of the 2nd USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 231–244.

SCHINDLER, J., GRIFFIN, J. L., LUMB, C. R., AND GANGER, G. R. 2002. Track-Aligned Extents: Matching Access Patterns to Disk Drive Characteristics. In *Proceedings of the 1st USENIX Conference on File and Storage Technologies*. USENIX Association, Monterey, CA, 259–274.

SCHINDLER, J., SCHLOSSER, S. W., SHAO, M., AND AILAMAKI, A. 2004. Atropos: A Disk Array Volume Manager for Orchestrated Use of Disks. In *Proceedings of the 3rd USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 159–172.

SCHLOSSER, S. W., SCHINDLER, J., PAPADOMANOLAKIS, S., SHAO, M., AILAMAKI, A., FALOUTSOS, C., AND GANGER, G. R. 2005. On multidimensional data and modern disks. In *Proceedings of the 4th USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 225–238.

SCHMIDT, A., WAAS, F., KERSTEN, M., FLORESCU, D., CAREY, M. J., MANOLESCU, I., AND BUSSE, R. 2001. Why and how to benchmark xml databases. *ACM SIGMOD Record 30,* 3 (September), 27–32.

SCHMUCK, F. AND HASKIN, R. 2002. GPFS: A Shared-Disk File System for Large Computing Clusters. In *Proceedings of the 1st USENIX Conference on File and Storage Technologies*. USENIX Association, Monterey, CA, 231–244.

SELTZER, M. I., GANGER, G. R., MCKUSICK, M. K., SMITH, K. A., SOULES, C. A. N., AND STEIN, C. A. 2000. Journaling Versus Soft Updates: Asynchronous Meta-data Protection in File Systems. In *Proc. of the Annual USENIX Technical Conference*. USENIX Association, San Diego, CA, 71–84.

SELTZER, M. I., KRINSKY, D., SMITH, K. A., AND ZHANG, X. 1999. The Case for Application-Specific Benchmarking. In *Proceedings of the IEEE Workshop on Hot Topics in Operating Systems (HOTOS)*. ACM, Rio Rica, AZ, 102–107.

SHEIN, B., CALLAHAN, M., AND WOODBURY, P. 1989. NFSSTONE: A network file server performance benchmark. In *Proceedings of the Summer USENIX Technical Conference*. USENIX Association, Baltimore, MD, 269–275.

SHEPLER, S. 2005. Nfs version 4. In *Proceedings of the Annual USENIX Technical Conference*. USENIX Association, Anaheim, CA. `http://mediacast.sun.com/share/shepler/20050414_usenix_ext.pdf`.

SHRIRA, L. AND XU, H. 2006. Thresher: An efficient storage manager for copy-on-write snapshots. In *Proceedings of the Annual USENIX Technical Conference*. USENIX Association, Boston, MA, 57–70.

SIVATHANU, G., SUNDARARAMAN, S., AND ZADOK, E. 2006. Type-Safe Disks. In *Proceedings of the 7th Symposium on Operating Systems Design and Implementation*. ACM SIGOPS, Seattle, WA, 15–28.

SIVATHANU, M., BAIRAVASUNDARAM, L. N., ARPACI-DUSSEAU, A. C., AND ARPACI-DUSSEAU, R. H. 2004. Life or Death at Block-Level. In *Proceedings of the 6th Symposium on Operating Systems Design and Implementation*. ACM SIGOPS, San Francisco, CA, 379–394.

SIVATHANU, M., BAIRAVASUNDARAM, L. N., ARPACI-DUSSEAU, A. C., AND ARPACI-DUSSEAU, R. H. 2005. Database-aware semantically-smart storage. In *Proceedings of the 4th USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 239–252.

SIVATHANU, M., PRABHAKARAN, V., ARPACI-DUSSEAU, A. C., AND ARPACI-DUSSEAU, R. H. 2004. Improving Storage System Availability with D-GRAID. In *Proceedings of the 3rd USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 15–30.

SIVATHANU, M., PRABHAKARAN, V., POPOVICI, F. I., DENEHY, T. E., ARPACI-DUSSEAU, A. C., AND ARPACI-DUSSEAU, R. H. 2003. Semantically-Smart Disk Systems. In *Proceedings of the 2nd USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 73–88.

SMALL, C., GHOSH, N., SALEEB, H., SELTZER, M., AND SMITH, K. 1997. Does Systems Research Measure Up? Tech. Rep. TR-16-97, Harvard University. November.

SMITH, K. A. AND SELTZER, M. I. 1997. File System Aging - Increasing the Relevance of File System Benchmarks. In *Proceedings of the 1997 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*. ACM SIGOPS, Seattle, WA, 203–213.

SNIA. 2007. SNIA - Storage Network Industry Association: IOTTA Repository. `http://iotta.snia.org`.

SOBTI, S., GARG, N., ZHANG, C., YU, X., KRISHNAMURTHY, A., AND WANG, R. 2002. PersonalRAID: Mobile Storage for Distributed and Disconnected Computers. In *Proceedings of the 1st USENIX Conference on File and Storage Technologies*. USENIX Association, Monterey, CA, 159–174.

SOULES, C. A. N., GOODSON, G. R., STRUNK, J. D., AND GANGER, G. R. 2003. Metadata Efficiency in Versioning File Systems. In *Proceedings of the 2nd USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 43–58.

SPADAVECCHIA, J. AND ZADOK, E. 2002. Enhancing NFS Cross-Administrative Domain Access. In *Proceedings of the Annual USENIX Technical Conference, FREENIX Track*. USENIX Association, Monterey, CA, 181–194.

SPC. 2007. Storage Performance Council. `www.storageperformance.org`.

SPEC. 2001. SPEC SFS97_R1 V3.0. `www.spec.org/sfs97r1`.

SPEC. 2003. SPEC SMT97. `www.spec.org/osg/smt97/`.

SPEC. 2004. SPEC SDM Suite. `www.spec.org/osg/sdm91/`.

SPEC. 2005a. SPECweb99. `www.spec.org/web99`.

SPEC. 2005b. The SPEC Organization. `www.spec.org/`.

SPEC. 2007. SPECviewperf 9. `www.spec.org/gpc/opc.static/vp9info.html`.

STEIN, C. A., HOWARD, J. H., AND SELTZER, M. I. 2001. Unifying file system protection. In *Proceedings of the Annual USENIX Technical Conference*. USENIX Association, Boston, MA, 79–90.

STRUNK, J. D., GOODSON, G. R., SCHEINHOLTZ, M. L., SOULES, C. A. N., AND GANGER, G. R. 2000. Self-Securing Storage: Protecting Data in Compromised Systems. In *Proceedings of the 4th Usenix Symposium on Operating System Design and Implementation*. USENIX Association, San Diego, CA, 165–180.

SWARTZ, K. L. 1996. The Brave Little Toaster Meets Usenet. In *Proceedings of the 10th USENIX System Administration Conference (LISA)*. USENIX Association, Chicago, IL, 161–170.

TAN, Y., WONG, T., STRUNK, J. D., AND GANGER, G. R. 2005. Comparison-based file server verification. In *Proceedings of the Annual USENIX Technical Conference*. USENIX Association, Anaheim, CA, 121–133.

TANG, D. 1995. Benchmarking Filesystems. Tech. Rep. TR-19-95, Harvard University.

TANG, D. AND SELTZER, M. 1994. Lies, Damned Lies, and File System Benchmarks. Tech. Rep. TR-34-94, Harvard University. December. In VINO: The 1994 Fall Harvest.

THERESKA, E., SCHINDLER, J., BUCY, J., SALMON, B., LUMB, C. R., AND GANGER, G. R. 2004. A Framework for Building Unobtrusive Disk Maintenance Applications. In *Proceedings of the 3rd USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 213–226.

TIAN, L., FENG, D., JIANG, H., ZHOU, K., ZENG, L., CHEN, J., WANG, Z., AND SONG, Z. 2007. Pro: A popularity-based multi-threaded reconstruction optimization for raid-structured storage systems. In *Proceedings of the 5th USENIX Conference on File and Storage Technologies*. USENIX Association, San Jose, CA, 277–290.

TOLIA, N., HARKES, J., KOZUCH, M., AND SATYANARAYANAN, M. 2004. Integrating Portable and Distributed Storage. In *Proceedings of the 3rd USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 227–238.

TOLIA, N., KOZUCH, M., SATYANARAYANAN, M., KARP, B., BRESSOUD, T., AND PERRIG, A. 2003. Opportunistic Use of Content Addressable Storage for Distributed File Systems. In *Proceedings of the Annual USENIX Technical Conference*. USENIX Association, San Antonio, TX, 127–140.

TRANSACTION PROCESSING PERFORMANCE COUNCIL. 2005. Transaction Processing Performance Council. `www.tpc.org`.

VEERARAGHAVAN, K., MYRICK, A., AND FLINN, J. 2007. Cobalt: Separating content distribution from authorization in distributed file systems. In *Proceedings of the 5th USENIX Conference on File and Storage Technologies*. USENIX Association, San Jose, CA, 231–244.

VERITAS SOFTWARE. 1999. VERITAS File Server Edition Performance Brief: A PostMark 1.11 Benchmark Comparison. Tech. rep., Veritas Software Corporation. June. `http://eval.veritas.com/webfiles/docs/fsedition-postmark.pdf`.

VERITEST. 2002. NetBench. `www.veritest.com/benchmarks/netbench/`.

VILAYANNUR, M., NATH, P., AND SIVASUBRAMANIAM, A. 2005. Providing tunable consistency for a parallel file store. In *Proceedings of the 4th USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 17–30.

VOGELS, W. 1999. File System Usage in Windows NT 4.0. In *Proceedings of the 17th ACM Symposium on Operating Systems Principles*. ACM, Charleston, SC, 93–109.

W. AKKERMAN. 2002. strace software home page. `www.liacs.nl/~wichert/strace/`.

WACHS, M., ABD-EL-MALEK, M., THERESKA, E., AND GANGER, G. R. 2007. Argon: Performance insulation for shared storage servers. In *Proceedings of the 5th USENIX Conference on File and Storage Technologies*. USENIX Association, San Jose, CA, 61–76.

WANG, A. A., REIHER, P., POPEK, G. J., AND KUENNING, G. H. 2002. Conquest: Better Performance Through A Disk/Persistent-RAM Hybrid File System. In *Proceedings of the Annual USENIX Technical Conference*. USENIX Association, Monterey, CA, 15–28.

WANG, R. Y., ANDERSON, T. E., AND PATTERSON, D. A. 1999. Virtual Log Based File Systems for a Programmable Disk. In *Proceedings of the 3rd Symposium on Operating Systems Design and Implementation*. ACM SIGOPS, New Orleans, LA, 29–44.

WANG, Y. AND MERCHANT, A. 2007. Proportional-share scheduling for distributed storage systems. In *Proceedings of the 5th USENIX Conference on File and Storage Technologies*. USENIX Association, San Jose, CA, 47–60.

WATSON, A. AND NELSON, B. 1992. LADDIS: A multi-vendor and vendor-neutral SPEC NFS benchmark. In *Proceedings of the 6th USENIX Systems Administration Conference (LISA VI)*. USENIX Association, Long Beach, CA, 17–32.

WEDDLE, C., OLDHAM, M., QIAN, J., WANG, A. A., REIHER, P., AND KUENNING, G. 2007. Paraid: A gear-shifting power-aware raid. In *Proceedings of the 5th USENIX Conference on File and Storage Technologies*. USENIX Association, San Jose, CA, 245–260.

WEIL, S., BRANDT, S., MILLER, E., LONG, D., AND MALTZAHN, C. 2006. Ceph: A Scalable, High-Performance Distributed File System. In *Proceedings of the 7th Symposium on Operating Systems Design and Implementation*. ACM SIGOPS, Seattle, WA, 307–320.

WITTLE, M. AND KEITH, B. E. 1993. LADDIS: The next generation in nfs file server benchmarking. In *Proceedings of the Summer USENIX Technical Conference*. USENIX Association, Cincinnati, OH, 111–128.

WRIGHT, C. P., DAVE, J., AND ZADOK, E. 2003. Cryptographic File Systems Performance: What You Don't Know Can Hurt You. In *Proceedings of the 2nd IEEE International Security In Storage Workshop*. IEEE Computer Society, Washington, DC, 47–61.

WRIGHT, C. P., JOUKOV, N., KULKARNI, D., MIRETSKIY, Y., AND ZADOK, E. 2005. Auto-pilot: A Platform for System Software Benchmarking. In *Proceedings of the Annual USENIX Technical Conference, FREENIX Track*. USENIX Association, Anaheim, CA, 175–187.

WRIGHT, C. P., MARTINO, M., AND ZADOK, E. 2003. NCryptfs: A Secure and Convenient Cryptographic File System. In *Proceedings of the Annual USENIX Technical Conference*. USENIX Association, San Antonio, TX, 197–210.

YU, X., GUM, B., CHEN, Y., WANG, R. Y., LI, K., KRISHNAMURTHY, A., AND ANDERSON, T. E. 2000. Trading Capacity For Performance. In *Proceedings of the 4th Usenix Symposium on Operating System Design and Implementation*. USENIX Association, San Diego, CA, 243–258.

YUMEREFENDI, A. R. AND CHASE, J. S. 2007. Strong accountability for network storage. In *Proceedings of the 5th USENIX Conference on File and Storage Technologies*. USENIX Association, San Jose, CA, 77–92.

ZADOK, E. 2002. Overhauling Amd for the '00s: A Case Study of GNU Autotools. In *Proceedings of the Annual USENIX Technical Conference, FREENIX Track*. USENIX Association, Monterey, CA, 287–297.

ZADOK, E., ANDERSON, J. M., BĂDULESCU, I., AND NIEH, J. 2001. Fast Indexing: Support for size-changing algorithms in stackable file systems. In *Proceedings of the Annual USENIX Technical Conference*. USENIX Association, Boston, MA, 289–304.

ZADOK, E., BĂDULESCU, I., AND SHENDER, A. 1999. Extending file systems using stackable templates. In *Proceedings of the Annual USENIX Technical Conference*. USENIX Association, Monterey, CA, 57–70.

ZADOK, E. AND NIEH, J. 2000. FiST: A Language for Stackable File Systems. In *Proc. of the Annual USENIX Technical Conference*. USENIX Association, San Diego, CA, 55–70.

ZHANG, C., YU, X., KRISHNAMURTHY, A., AND WANG, R. Y. 2002. Configuring and Scheduling an Eager-Writing Disk Array for a Transaction Processing Workload. In *Proceedings of the 1st USENIX Conference on File and Storage Technologies*. USENIX Association, Monterey, CA, 289–304.

ZHANG, Z. AND GHOSE, K. 2003. yFS: A Journaling File System Design for Handling Large Data Sets with Reduced Seeking. In *Proceedings of the 2nd USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 59–72.

ZHOU, Y., PHILBIN, J., AND LI, K. 2001. The multi-queue replacement algorithm for second level buffer caches. In *Proceedings of the Annual USENIX Technical Conference*. USENIX Association, Boston, MA, 91–104.

ZHU, N., CHEN, J., AND CHIUEH, T. 2005. Tbbt: Scalable and accurate trace replay for file server evaluation. In *Proceedings of the 4th USENIX Conference on File and Storage Technologies*. USENIX Association, San Francisco, CA, 323–336.

ZHU, Q., CHEN, Z., TAN, L., ZHOU, Y., KEETON, K., AND WILKES, J. 2005. Hibernator: Helping Disk Arrays Sleep through the Winter. In *Proceedings of the 20th ACM Symposium on Operating Systems Principles*. Brighton, UK, 177–190.