

Thinking about benchmarks

A benchmark's primary purpose is to serve as a standard for comparison. By running the same workload with the same initial conditions on the same hardware in the same execution environment, researchers/practitioners can fairly compare two competing system designs (or different parameter choices for the same system design). *Proper* benchmarking is an important part of the scientific process.

In previous courses, we've evaluated algorithm performance in terms of best case, worst case, average case, and expected case. What a "typical systems user" in the wild probably cares most about is the representative case (although some companies care very much about the worst case). A big challenge in the storage community is creating a so-called representative state before conducting a system evaluation. A realistic workload is significantly less meaningful if it is run under unrealistic conditions.

Thought questions

For these questions, consider the context of a local file system (e.g., your ext4 or similar file system on your Linux laptop) and think about contents on disk, memory, and any active work that being performed on the machine.

1. What is an appropriate *initial state* on which to run a benchmark?
2. What is the *steady state* of your file system?
3. Think back to the file system designs we've covered. Do you think that the initial/steady state should be modeled as a *path* or a *point*, and why?
4. How hard is it to create an exact representation of your current laptop state? What are some of the most important aspects to recreate?

File System Aging---Increasing the Relevance of File System Benchmarks (Smith and Seltzer)

Motivation and Goals. Measuring the performance of an empty file system is (arguably) poor science. Creating a realistically aged file system is hard but necessary work if you want to make representative measurements. Smith and Seltzer create *aging tools* that capture some key aspects of real file system workloads; these tools can then be used to replay important features of those workloads to bring a system into a reproducible state that exhibits representative aging.

Methods. The authors collected *file system snapshots* over several years. They then simulate file system workloads based on the path that those snapshots describe; they also simulate reasonable operations that, by nature, cannot be captured in the snapshot state.

- **Ideally:** Start from an empty disk and faithfully replay a trace of real operations that occur over months/years, bringing that file system to an almost full state.
- **Challenge:** Collecting a perfect trace is expensive and impractical.
- **Solution:**
 - i. Collect a set of file system *snapshots*.
 - ii. Given a sequence of snapshots, generate a plausible aging workload (series of file system operations, e.g., creates, deletes, modifications, etc.) that brings the file system state of one snapshot in the sequence to the next snapshot's file system state.
 - iii. In addition, also simulate the creation/deletion of short-lived files that do not appear in any of the snapshots, but may exist in the transient states between two snapshots.

They authors also verify that the workloads faithfully mirror the snapshot states. In this way, they show that their tool does in fact create *representative* aged file system states.

New term: layout score. The layout score is used to quantify the amount of file system fragmentation.

- An individual file's layout score is the fraction of its blocks that are optimally located within the LBA space (e.g., perfectly contiguous).
 - 1 is perfect, 0 is the worst possible score.
 - The first block in a file is not counted.
- The aggregate layout score of all the files in a file system measures its overall fragmentation.

Thought questions

1. How does a file's metadata placement affect a file's layout score? Does metadata placement affect file system performance (think back to the FFS motivation/discussion in class)?
2. The layout score is for individual files. Does it capture the relationship between files? How about related files?
3. What aspects of a file system workload cannot be inferred from a sequence of snapshots? How do the authors attempt to simulate those aspects, and is there anything that you think they are still missing?

How to Fragment Your File System (Conway et al.)

To continue to be representative as the way we interact with our storage evolves (think about smart phones usage patterns vs. 1990s PCs...), we would need to recreate Smith and Seltzer's approach: our tools would need to reflect new aging patterns induced by new patterns of file system usage. Snapshot collection and tool validation is expensive.

The layout score, while informative, also has some limitations. It only captures the fragmentation of an individual file's data. We also care about metadata placement and the placement of files that are related to each other.

Conway et al. creates a new tool based on the Git version control system that (1) generates an aged file system and (2) measures the file system's aging progression.

Fragmentation categories Conway et al. describes different categories of fragmentation that affect the performance of typical file system workloads. Their measurement of a file system's layout is based on performing *recursive grep* operations, where a file system tree is recursively traversed and each block of each file is read (both data and metadata) in breadth-first search order.

- intrafile fragmentation involves blocks from the same file (e.g., as measured by Smith and Selzer's layout score)
- interfile fragmentation involves blocks from two different files (e.g., two adjacent files in a directory)
- metadata fragmentation involves at least one metadata block (e.g., the distance between a block in a file's inode and the data block it describes OR the distance between two inodes describing files in the same directory)

Methods. The aging tool is simple works as follows:

1. Successively checkout commits in a git repository, starting from the first.
2. After every 100 git pulls, measure the performance of a recursive grep of the file system (normalized by seconds / GiB to reflect changes in the file system size)
3. Compare the aged performance against a "clean" file system where the data is copied to a fresh partition. In the "clean" file system, the file system can ideally place files without confounding updates.

Notes: When measuring a system, it is important to *isolate* the target effects. Conway et al. wants to isolate the effects of aging, so they create an unaged baseline, and they normalize their measurement (grep cost in seconds/GiB) to remove the effects of changing file system sizes. They also perform *cold cache* measurements so that all of the costs of seeking are paid by the system.

They later ask targeted questions that they answer by again isolating each of the effects.

1. Do SSDs (which have different random access behaviors than HDDs) eliminate the effects of aging? They run the same experiments on an SSD.
2. Does caching eliminate the effects of aging (and if so, what is the sensitivity to different cache sizes)? They perform warm-cache tests with different amounts of memory to identify trends and see how large the effect is.
3. Does increasing disk capacity eliminate the effects of aging (with the argument being that file systems have significantly more choice for their allocations when there is excess capacity)? Using the same physical disk, they perform the experiments with different partition sizes (which effectively expands the LBA space and eliminates disk geometry effects).

The answer to all of these questions is no. Aging is a problem as long as your working set exceeds cache.

Thought Questions

1. Is recursive grep actually a good workload? What are some of its limitations? Its strengths?
2. What operations might be missed by using "git pull" as a workload? How might you modify the tool to incorporate those operations?
3. Is the git replay workload representative of your file system usage? What does it actually capture?

More info

The Conway et al. paper was a magazine article written to distill a larger paper down to a bite-size format. For more details, the full workshop paper (and talk) is available online in the FAST 2017 proceedings. The full paper is **File Systems Fated for Senescence? Nonsense, Says Science!** by Alex Conway, Ainesh Bakshi, Yizheng Jiao, Yang Zhan, Michael A. Bender, William Jannen, Rob Johnson, Bradley C. Kuszmaul, Donald E. Porter, Jun Yuan, and Martin Farach-Colton.

Generating Realistic Impressions for File-System Benchmarking by Nitin Agrawal, Andrea C. Arpaci-Dusseau and Remzi H. Arpaci-Dusseau is another important aging paper. That paper introduces a tool to synthetically age a file system (meaning the tool doesn't replay a real workload, but instead simulates a workload based on a small number of parameters). The tool has many advantages:

- A workload can be compactly encoded and shared
- There are many parameters that let you control the settings
- Reproducibility is easy