

## 1 Summary

This paper introduced the novel log-structured file system: a new technique for disk storage management (at the time). Unlike traditional disk storage management, this new system writes modifications to disk sequentially (like a log). After highlighting the need for this sort of improvement, the authors demonstrate how their implementation, Sprite LFS, outperforms rival Unix file systems by an order of magnitude for small-file writes, while matching/exceeding Unix performance for reads and large writes. Moreover, Sprite LFS can use up to 7x - 14x more disk bandwidth when the overhead for cleaning is included compared to Unix file systems.

## 2 Strengths of the paper

The authors made a concise, yet strong, point in §2.3 about the problems faced by existing systems like UNIX to portray the need for their log-based file system. Findings like inefficient spreading of information resulting in excessive small accesses and synchronous writes that cause the application to wait until the write is complete are sufficient to convince the technical reader.

The key ideas of the *segment usage table* as the data structure used in the Sprite LFS to support the cost-benefit cleaning policy was well-explained to readers who have no knowledge of this data structure's existence. Previous sections related to the needs of the system show why the segment usage table is a well-designed data structure for this implementation.

The paper is organized in a very easy-to-follow manner. The authors state their goals and the motivation behind them, describe the system they developed to meet these goals, and finally, discuss their results based on experimentation while comparing with existing systems. What else is expected of such a paper?

## 3 Weakness of the paper

I'm being nitpicky here because I can't think of anything else to say. But this is the first time I've seen a related work section right before the conclusion. Typically, papers tend to put this section right around the introduction or literature review earlier on, so readers are aware of the context this paper aims to operate within. I'm not too sure what the authors were aiming to achieve by placing this section where they did.

## 4 Future work opportunities

It's been some time since this paper was published, and more advanced ideas leveraging this system have since been implemented. I wonder in what ways these newer systems decided to retain the features they did and do without the ones they do without.