

CS457/657 Database Management Systems

Paper Review 2

Graduate students taking CS657 are required to select one of the following two papers to read and write a full review. This is an individual assignment, no collaboration allowed.

1. The Hadoop Distributed File System, <https://dl.acm.org/citation.cfm?id=1914427>
2. Ceph: a scalable, high-performance distributed file system,
<https://dl.acm.org/citation.cfm?id=1298485>

Here are some suggestions on how to review a research paper:

1. <https://www.cse.unr.edu/~dfz/download/p83-keshavA.pdf>
2. <https://homes.cs.washington.edu/~mernst/advice/review-technical-paper.html>

There is no required format for the review. But at the very least, we expect the following questions to be answered in your review:

- Is the presentation clear? Or, would you present your research work this way? Why, or why not?
- What problems does the paper try to solve?
- Are the problems important problems? If so, why?
- Are the problems hard problems? Why, or why not?
- What were the state-of-the-art techniques before the paper got published?
- Is the proposed technique proposed in the paper better than the state-of-the-art? If so, by how much?
- At a very high level, why does the proposed technique work?
- How did the authors convince the readers that the proposed techniques really work?
- What are the limitations of the proposed work?
- On what aspects do you think the paper and its proposed techniques can be improved?

We don't require a minimal length of words/pages for the review either. But if you really understand the paper, you should be able to elaborate your answer to each of the questions instead of, for instance, a single sentence of a few words.