

Reading List and Instructions

This is the reading list for the seminar. Each student should choose one paper/topic.

Timeline for paper and presentation assignment

- List published by **16.03.2025**; I will go over the list in the following class and will say a few words about each topic.
- Presentations will start around Week 7, namely **08.05.2025**
- Each student should send me as soon as possible, but no later than **30.03.2025** midnight, an email with his/her preferences for:
 - at least 2 choices papers/topics
 - 2-3 choices for a date
 - Possible dates:
 - **08.05.25**
 - **15.05.25**
 - **22.05.25**
 - **29.05.25**
 - **05.06.25**
 - **12.05.25**
- If you need to consult with me in order to make the choice - I will be available after class, or send me an email and we will schedule time to talk on Zoom.
- Final paper assignment and date for presentation – will be published on **03.04.2025**
- Note:
 - I will do the assignments based on your preferences
 - I will do my best to accommodate everyone as much as possible (in terms of paper choices and dates) but eventually you have to accept the assignment.
- Once the timeline is set, every student should arrange 45 mins with me (possibly on zoom) two weeks before the talk to go over the slides. Note that we may need another session.

How to choose a paper?

Note: Some of the papers need an OpenAccess to the ACM/IEEE Digital Library. You can access them through one of the links [here](#).

If you still have a problem accessing the papers please let me know.

- Read the abstracts of all the papers and select those that seem relevant to you
- Read “about” the topic on Wikipedia or relate blogs to get a feel to the topic
- For the relevant ones, look at the paper
 - o Read the introduction & background, the conclusions, the headings, look at the figures and graphs and their captions
- Try to pick a topic that interests you. If you have prior knowledge about the topic (from work for example) it can help.

(You can read “[Efficient Reading of Papers in Science and Technology](#)”)

Preparing your presentation

Read the paper and make sure you get the overall idea. Then read it again more deeply, take notes, challenge what is said in the paper until you understand it fully. Only then start answering those questions and this will guide you thru the presentation.

Also collect some additional information about the topic covered in the paper. Related sources:

- Wikipedia
- Blogs
- Keynote presentations in conferences

The presentation should have the following components:

- Details about the paper (authors & affiliations, when and where it was published, general area/topic)
- Presented by...
- Setting the stage: Motivation, background, assumptions
- Problem addressed
- Previous approaches
- In a nutshell: what’s the main contribution of the paper
- Detailed section:
 - o The main solution: algorithm, systems, method
- Analysis
 - o General analysis
 - o Measurements: What were the experiments, the workloads
 - o Systems characteristics
 - o Pick one or two graphs and analyze them thoroughly
- Summary
- Your opinion and reflections on the paper
- How influential was the paper and the system, e.g.:
 - o Citation count
 - o Blogs, articles that refer to this paper or system
 - o Active implementations

Be prepared to give a 50 minutes presentation, following by Questions & Answers

Make sure not to “cut and paste” slides from publicly available sources. I don’t mind that you use some pictures, figures etc. but the presentation should be YOURS, and you should think about the logic, the organization, and the content, and what you really want to say.

Using ChatGPT or related LLM services to help with the preparation is OK, but be smart about it. Apply your judgement and make sure you understand what you present! You should be prepared for hard questions.

Proposed Reading List

I am also open to suggestions for other papers.

Next to each paper we mark the topic(s) addressed by the paper.

- [DS] Distribute Systems, consensus
- [CS] Cloud Storage, distributed storage, file systems
- [BDF] BigData processing frameworks
- [A] BigData Analytics support

1. Dynamo (Amazon) and Spanner (Google) [DS, CS]

Giuseppe DeCandia, Deniz Hastorun, Madan Jampani, et al.: “[Dynamo: Amazon’s Highly Available Key-Value Store](#),” at *21st ACM Symposium on Operating Systems Principles (SOSP)*, October 2007.

James C. Corbett et al. Spanner: Google’s Globally-Distributed Database. *10th USENIX Symposium on Operating Systems Design and Implementation (OSDI)* 2012.

<https://www.usenix.org/system/files/conference/osdi12/osdi12-final-16.pdf>

2. BigTable (Google) * [CS]

Chang, F., Dean, J., Ghemawat, S., Hsieh, W. C., Wallach, D. A., Burrows, M., ... & Gruber, R. E. (2008). Bigtable: A distributed storage system for structured data. *ACM Transactions on Computer Systems (TOCS)*, 26(2), 1-26. Also in *OSDI* 2006,

<https://www.usenix.org/legacy/event/osdi06/tech/chang/chang.pdf>

3. Local Reconstruction Codes in Azure OS (Microsoft) * [CS]

Huang, C., Simitci, H., Xu, Y., Ogus, A., Calder, B., Gopalan, P., ... & Yekhanin, S. (2012). Erasure coding in windows azure storage. In *2012 USENIX Annual Technical Conference (USENIX ATC 12)* (pp. 15-26). (Best paper)

https://www.usenix.org/system/files/conference/atc12/atc12-final181_0.pdf

4. Consensus and Coordination in large scale systems – Zookeeper (Yahoo!) and Raft [DS]

4.1 Zookeeper

Hunt, P., Konar, M., Junqueira, F. P., & Reed, B. (2010). {ZooKeeper}: Wait-free Coordination for Internet-scale Systems. In *2010 USENIX Annual Technical Conference (USENIX ATC 10)*.

https://www.usenix.org/legacy/event/atc10/tech/full_papers/Hunt.pdf

4.2 Raft

Diego Ongaro and John Ousterhout (2014). In Search of an Understandable Consensus Algorithm. USENIX ATC 2014.

<https://www.usenix.org/system/files/conference/atc14/atc14-paper-ongaro.pdf>

5. Cloud file systems: GFS (Google) and HDFS (Yahoo!) * [CS]

[SOSP 2003]

Ghemawat, S., Gobioff, H., & Leung, S. T. (2003, October). The Google file system. In *Proceedings of the nineteenth ACM symposium on Operating systems principles* (pp. 29-43).

<https://dl.acm.org/doi/pdf/10.1145/945445.945450?download=true>

[MSST 2010]

Shvachko, K., Kuang, H., Radia, S., & Chansler, R. (2010, May). The hadoop distributed file system. In *2010 IEEE 26th symposium on mass storage systems and technologies (MSST)* (pp. 1-10). IEEE.

<https://storageconference.us/2010/Papers/MSST/Shvachko.pdf>

6. Batch processing: MapReduce (Google) and Spark (Berkeley, AMP Lab) [BDF, CS]

Here talk about MapReduce vs. Spark in general, and choose one of the papers in depth.

6.1 MapReduce

Jeffrey Dean, Sanjay Ghemawat. MapReduce: Simplified Data Processing on Large Clusters. OSDI 2004: 137-150

https://www.usenix.org/legacy/publications/library/proceedings/osdi04/tech/full_papers/dean/dean.pdf

Also, see the CACM paper 4 years later

Jeffrey Dean and Sanjay Ghemawat. 2008. MapReduce: simplified data processing on large clusters. *Commun. ACM* 51, 1 (January 2008), 107–113.

<https://dl.acm.org/doi/pdf/10.1145/1327452.1327492>

6.2 Spark

Zaharia, M., Chowdhury, M., Das, T., Dave, A., Ma, J., McCauly, M., ... & Stoica, I. (2012). Resilient distributed datasets: A Fault-Tolerant abstraction for In-Memory cluster computing. In *9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12)* (pp. 15-28). (Best paper)

<https://www.usenix.org/system/files/conference/nsdi12/nsdi12-final138.pdf>

7. Security/Data Privacy (Microsoft) [A]

McSherry, F. D. (2009, June). Privacy integrated queries: an extensible platform for privacy-preserving data analysis. In *Proceedings of the 2009 ACM SIGMOD International Conference on Management of data* (pp. 19-30).

<https://css.csail.mit.edu/6.5660/2024/readings/pinq.pdf>

8. Data Streaming – Kafka By Linkedin

Kafka: a Distributed Messaging System for Log Processing. Kreps, Jay, Neha Narkhede, and Jun Rao NetDB 2011.

<https://pages.cs.wisc.edu/~akella/CS744/F17/838-CloudPapers/Kafka.pdf>

This is a highly influential paper. I expect that you cover more than the paper

9. Sky computing

Ion Stoica and Scott Shenker (2021). From cloud computing to sky computing. *Proceedings of the Workshop on Hot Topics in Operating Systems (HotOS '21)*.

<https://sigops.org/s/conferences/hotos/2021/papers/hotos21-s02-stoica.pdf>

Zongheng Yang et al. Shenker, Stoica (2023). SkyPilot: An Intercloud Broker for Sky Computing. Symposium on Networked Systems Design and Implementation (NSDI) 2023.

<https://www.usenix.org/system/files/nsdi23-yang-zongheng.pdf>

Other sources:

Keynote at Usenix ATC 2023 <https://www.usenix.org/conference/atc23/presentation/joint-keynote>

10. DataLake (and Lakehouse storage systems)

Michael Armbrust et al. and Matei Zaharia. 2020.

Delta lake: high-performance ACID table storage over cloud object stores. Proc. VLDB Endow. 13, 12 (August 2020), 3411–3424. <https://doi.org/10.14778/3415478.3415560>

Paras Jain, Peter Kraft, Conor Power, Tathagata Das, Ion Stoica and Matei Zaharia. Analyzing and Comparing Lakehouse Storage Systems. *Conference on Innovative Data Systems Research (2023)*.

<https://www.cidrdb.org/cidr2023/papers/p92-jain.pdf>