COMS 4113 Homeworks

Fall 2024

Anirudh Varanasi

Agenda

- Deadlines & How HWs factor into your grade
- Difficulty
  - Tips
  - Testing your assignments
- Overview of homeworks
  - Topics
- Grading Specifics
  - HW2 - HW5
- HW1 - in depth
- Grading for HW1
- Questions

Deadlines & How HWs factor into your grade

8 Homeworks in total
(excludes HW0)

- Best 6 out of 7 HWs will be used for HW1 through HW 4b
  - 60%
- HW5 is a must
  - 10%
- Total = 70% of your final grade
- 10% extra credit for perfect score on all 8 assignments

Homework	Submission Deadline	Weights
HW0	Mon, Sept 9 (2 days!)	0 (but required)
HW1	Mon, Sept 16 (1 week)	10%
HW2a	Mon, Sept 30 (2 weeks)	10%
HW2b	Mon, Oct 7 (1 week)	10%
HW3a	Mon, Oct 21 (2 weeks)	10%
HW3b	Mon, Oct 28 (1 week)	10%
HW4a	Mon, Nov 11 (2 weeks)	10%
HW4b	Mon, Nov 25 (1 week + 3 days)	10%
HW5	Mon, Dec 16 (3 weeks)	10%

Difficulty

- HW1 < HW2 < HW3 < HW4 ≈ HW5
- Part a < Part b

Tips across all homework assignments

- Read papers and understand the protocol before coding
- Frequently print your results when debugging a distributed system
- Start your part B before the deadline of part A
- When designing solution for part A, keep part B in mind
- Read ahead!
- Get familiar with what the tests expect for each homework assignment

Testing your assignments

- Run unit tests at least 100 times, maybe on different machines
- The result is not deterministic (especially for HW2-4), due to goroutine/thread scheduling
- Passing tests 50 times does not mean your code is correct, and your code may still fail on our grading machine!
  - You will be given GCP credits and instructions to setup your own instances
- We won’t add hidden unit tests
- You can also add print statements in the unit tests, just remove them before submitting -

Overview of Homeworks

Homework	Project	Related Topics
HW1	MapReduce	MapReduce, RPC
HW2	Primary/Backup Key/Value Datastore	Fault Tolerance
HW3	Paxos & Key/Value Datastore	Consensus, Paxos, Availability
HW4	Sharded Key/Value Datastore	Scalability, Paxos, Atomic commitment
HW5	Model Checking the Paxos Protocol	Testing & Model Checking

Themes common across all HWs

- TDD - the tests are provided for you
- A lot of the API skeletons are provided
  - You fill them out
  - You can add more
- Design, design, design
- Data structures
  - Some are provided
  - Some you modify
  - Some you can make yourself

Grading

- Unit tests are used to grade your assignments.
- Unit tests in the same homework each have the same weight
- Each unit tests will be run 50 times.
- Every time a unit test fails, the score of this unit test will be multiplied by 0.9.

#fails	0	1	2	3	4	5	6	7	8	25
score	100%	90%	81%	73%	66%	59%	53%	48%	43%	7%

HW1 - MapReduce

- Getting used to RPC & Go
- A simple Map/Reduce implementation
  - Part 1: Implement Map & Reduce
  - Part 2: Distribute Workload
  - Part 3: Account for workers failing

HW1 - MapReduce

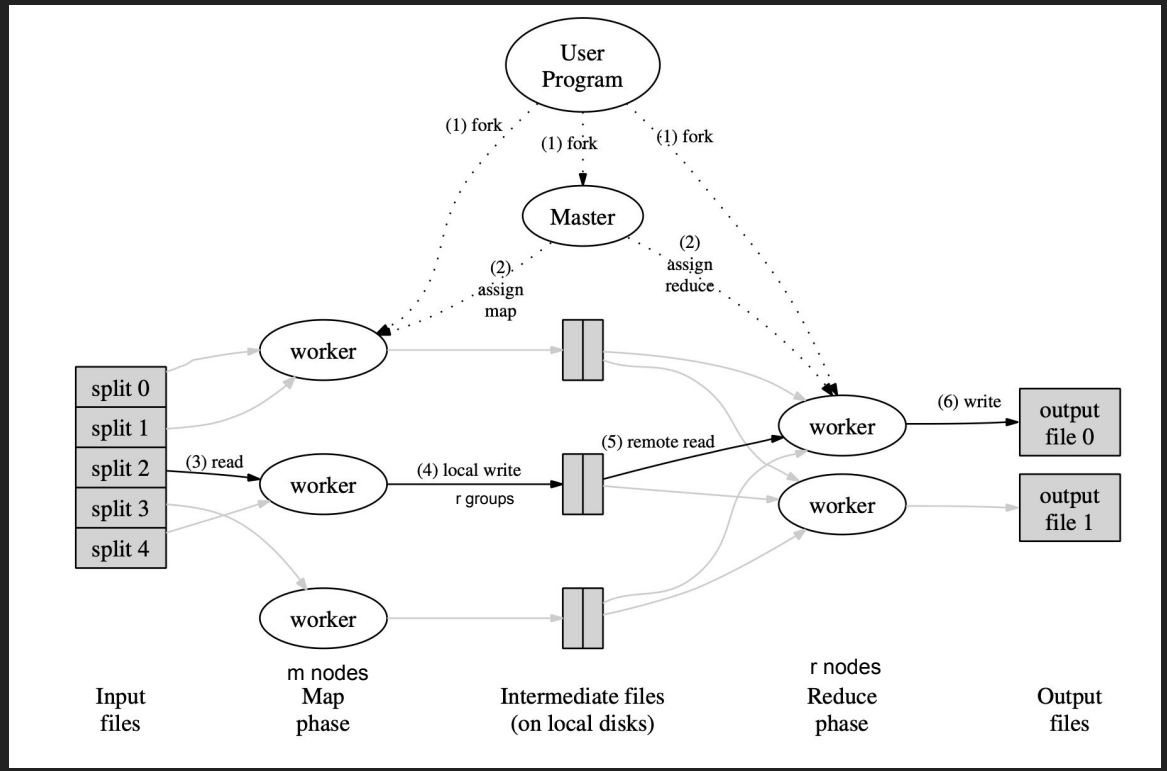

From the MapReduce paper (link)

HW1 - MapReduce - more in depth (WordCount)

- Split file: input-1.txt, input-2.txt, ..., input-m.txt
- Map Phase
  - *"..., a dog, a cat ..."* =>
  - `{"a": ["1", "1"], "dog": ["1"], "cat": ["1"], ...}`
- Intermediate files
  - int-file-1-1.txt: `{"a": ["1", "1"], "cat": ["1"], ...}`
  - int-file-1-2.txt: `{"dog": ["1"]...}`
  - …
  - int-file-m-r.txt
- Reduce Phase
  - `{"a": ["1", "1"], "cat": ["1"], ...},`
    `{"a": ["1", "1", "1"], "apple": ["1", "1"]}`=>
  - `{"a": "5"}, {"apple": "2"}, {"cat": "1"}, …`
- Output: output-1.txt, output-2.txt, ..., output-r.txt

HW1 - MapReduce - more in depth (con’t)

Part 1

- Straightforward implementation
  - Side-note: there a few ways to run this (sequential, etc).
  - The comments provide some instructions
- Map each word; Count the mappings

Part 2/3

- Map & Reduce are written for you
- You have to distribute the workload
  - Write a master that coordinates Map & Reduce jobs between workers
- Account for worker failures

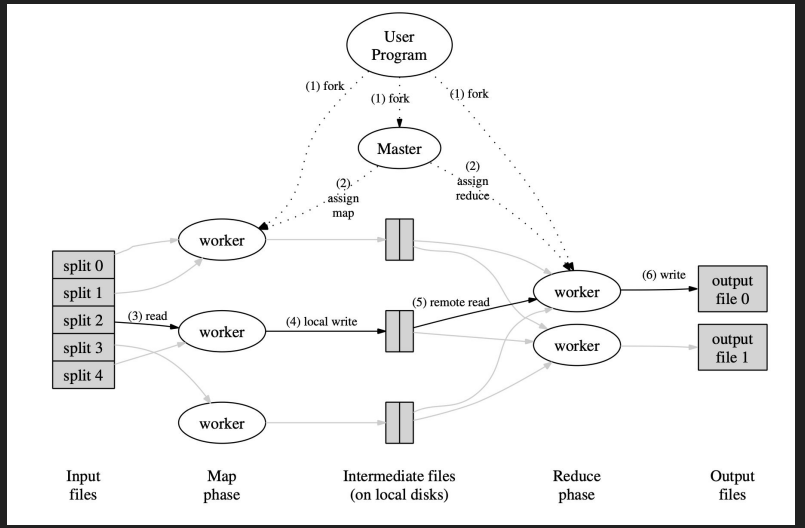

Grading for HW1

- Test 1 (20 points) (Part 1)
  - test-wc.sh under src/main.
  - If it passed, you'll receive 20 points, or for this part, the grade is 0
- Test 2 (100 points) (Part 2/3)
  - Follows aforementioned grading

The final grade (120 points) is scaled to 100.

For example, if you get

- 20 points in Test 1
- 70 points in Test 2
- Then you get (20+70)*⅚ = 75 for HW1

Special thanks to Juan Gutierrez for making most of these slides and Sida Huang who was the source of the content for some of these slides.