UNIVERSITY OF TARTU
Institute of Computer Science

LTAT.06.007 Distributed Systems

Practical Seminar 10

Farooq Ayoub

TEACHING ASSISTANT

Tartu, Estonia 20/04/2021

Recap

- **Check whether a given execution steps are Sequentially Consistent or Not**
  - Sequential Consistency
  - Check for Sequential Consistency and if not rearrange the Operations to make them Sequentially Consistent.

Agenda

- **Goal: Quorum-Based Protocols**
- **Content:**
  - Gifford's Scheme of Voting for supporting replicated writes
  - Apply Gifford Quorum-Based protocol to a set of examples and look for Conflicts.
- **Quiz**

Session Content

Description

- Basic idea about Quorum Based Protocols
- Gifford's Scheme
- Constraints followed in Gifford's Scheme
- Some Examples about Quorum

Observation

Instructions to complete this practical session can be found in the course website: https://https://courses.cs.ut.ee/2021/ds/spring/Main/Instructions6

Quorum-Based Protocols

- Resolves write-write or read-write conflicts
- Client processes are required to request and acquire the permission of multiple servers before reading and writing a replicated data item.
- Example protocol (on a distributed file system):
  - A process that wants to update a replicated file first contacts at majority of servers and get them to agree to do the update.
  - Once they agreed, the file is changed and a new version number is associated with the file.
  - To read a replicated file, a client must also contact at least half the servers plus one and ask them to send the version numbers associated with the file.
  - If all version numbers agree then the file is the most recent one.

Gifford's Scheme

- N replicas exist.
- To read a file, a client needs to assemble a Read Quorum:
  - An arbitrary collection of any NR servers.
- To write a file, a client needs to assemble a Write Quorum:
  - An arbitrary collection of any NW servers

Gifford’s Scheme - Constraints

The values NR and NW are subject to following constraints:

- NR + NW > N
- NW > N/2

The first constraint prevents Read-Write conflicts

The second constraint prevents Write-Write conflicts

Example 1

A correct choice of read and write set

- N=12, NR=3, NW = 10
- Both constraints are followed

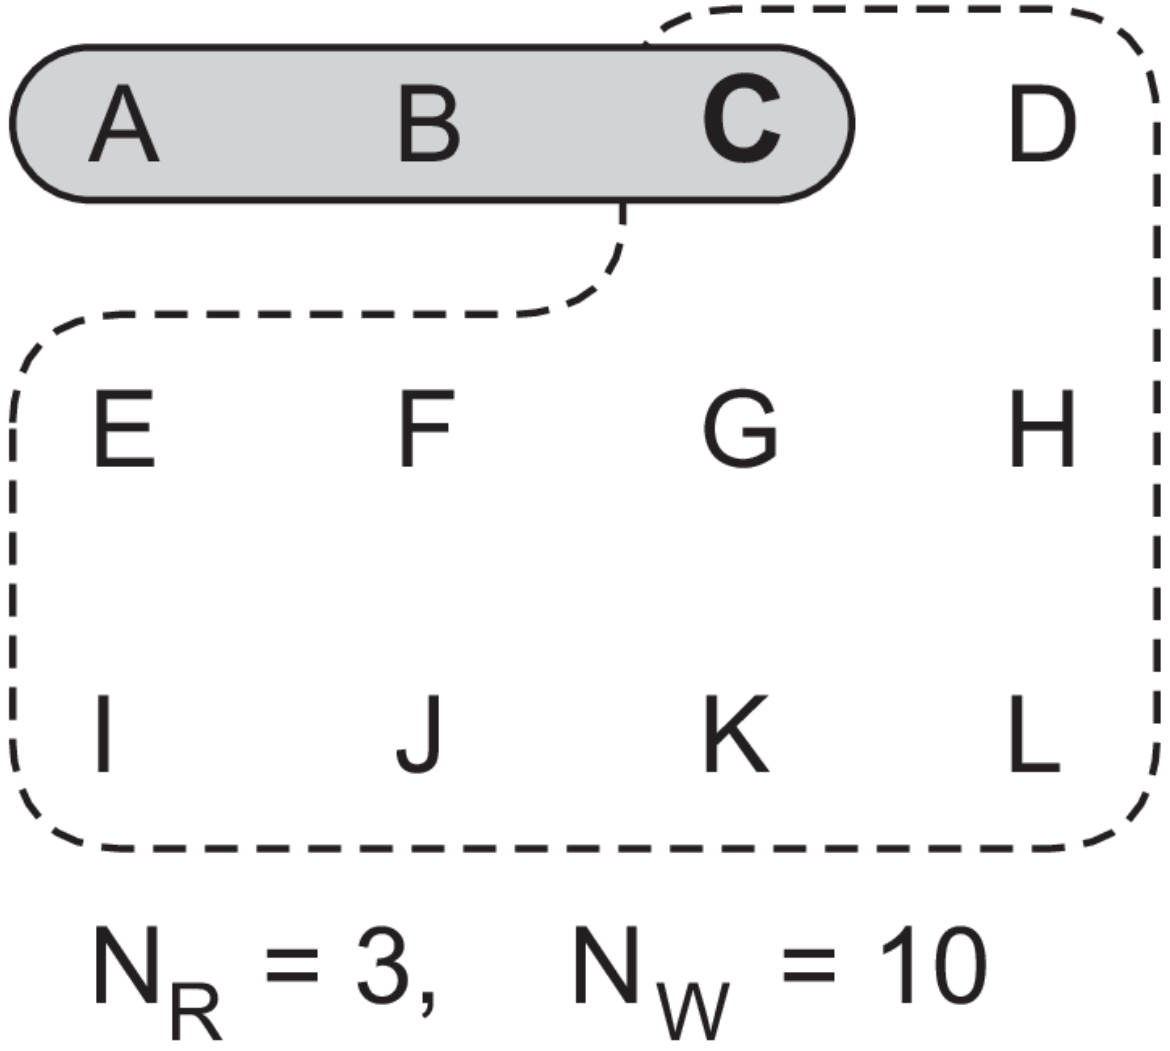

Example 2

- A Write-Write conflict may occur because NW <= N/2.
- One client may choose (A, B, C, E, F,G) as its write set and another client may choose (D,H, I, J,K, L) as its write set.
- Trouble as the two updates will both be accepted without detecting that they actually conflict.

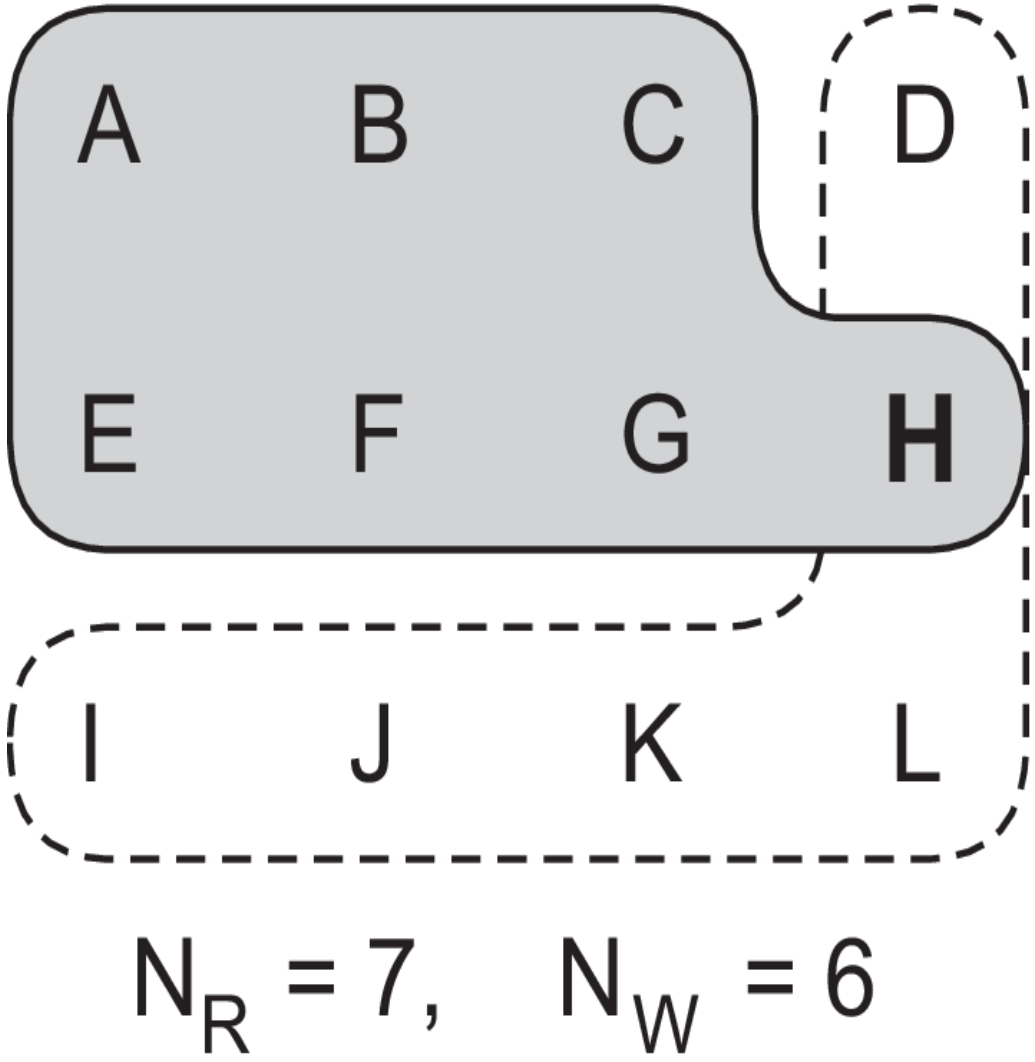

Example 3

- Sets NR to one, making it possible to read a replicated file by finding any copy and using it.
- A Special Case known as “Read-One, Write-All (ROWA)”.

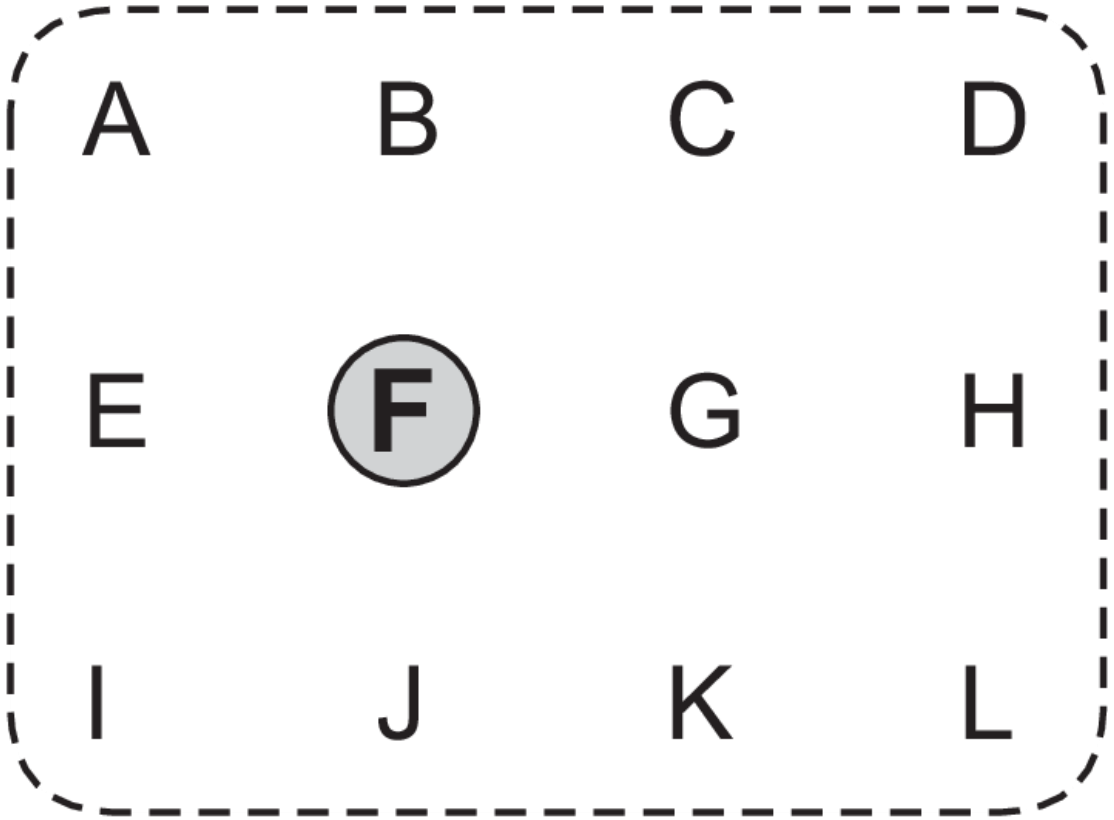

$N_R = 1, \quad N_W = 12$

Session Instructions at Course Page

Quiz

Content

- Lecture 11 (Consistency and Replication)
- Two attempts
  - One in Seminar Session
  - Next available until Monday 23:59 (Deadline)
- Open Quiz in Moodle
- Total Quiz Points = 100

Observation

Quiz review is available after the quiz is closed

UNIVERSITY OF TARTU
Institute of Computer Science

Questions?

E-mail: farooq.ayoub.dar@ut.ee