

INFO8002

Large-Scale Data Systems

Exercise Session #3

Academic year 2021-2022



Reminder

REMINDER :

Consensus

- **Consensus** mechanisms allows a distributed system to maintain some common state or to agree on some future actions.

- **FLP Theorem**

*“In an asynchronous network where messages may be delayed but not lost, there is **no consensus algorithm that is guaranteed to terminate** in every execution for all starting conditions, if at least one node may experience failure.”*

REMINDER :

Consensus VS Uniform Consensus

- Two types of consensus algorithms, differing on the guarantees they provide in the presence of faulty processes:

Module:

Name: *Consensus*, instance *c*.

Events:

Request: $\langle c, \text{Propose} \mid v \rangle$: Propose value v for consensus.

Indication: $\langle c, \text{Decide} \mid v \rangle$: Outputs a decided value v of consensus.

Properties:

C1: Termination: “Every correct process eventually decides some value.”

C2: Validity: “If a process decides v , then, v was proposed by some process.”

C3: Integrity: “No process decides twice.”

C4: Agreement: “No two correct process decide differently.”

Module:

Name: *UniformConsensus*, instance *uc*.

Events:

Request: $\langle uc, \text{Propose} \mid v \rangle$: Propose value v for consensus.

Indication: $\langle uc, \text{Decide} \mid v \rangle$: Outputs a decided value v of consensus.

Properties:

C1: Termination: “Every correct process eventually decides some value.”

C2: Validity: “If a process decides v , then, v was proposed by some process.”

C3: Integrity: “No process decides twice.”

C4: Uniform Agreement: “No two processes decide differently.”

REMINDER :

Consensus VS Uniform Consensus

- **Consensus** can be implemented in synchronous and partially synchronous systems:

Consensus:

- **Fail-stop - Hierarchical**
 - Broadcast: **Best effort**
 - Failure Detector : **Perfect**


Uniform Consensus:

- **Fail-stop - Hierarchical**
 - Broadcast: **Reliable & Best effort**
 - Failure Detector : **Perfect**
- **Fail-noisy - Leader Driven**
 - Use of Epoch Consensus using Eventual Leader Election.

PROBLEM 1

Version Control System


You are responsible for designing a system allowing decentralized version control.




Specify an architecture for this distributed version control system and **provide** a pseudo-implementation using consensus.

Version Control System

Write query



Read query



Version Control System

Problems:

1. How can we ensure that operations terminates ?

Version Control System

Problems:

1. **How can we ensure that operations terminates ?**

Property n°1: “Termination”

“If a correct process invokes an operation, then the operation eventually completes.”

2. **How can we ensure that commits are applied consistently on all remote clients?**

Version Control System

Problems:

1. How can we ensure that operations terminates ?

Property n°1: “Termination”

“If a correct process invokes an operation, then the operation eventually completes.”

2. How can we ensure that commits are applied consistently on all remote clients?

Property n°2: “Agreement”

“All correct processes applies the same sequence of operation.”

Version Control System

Module Specification:

Module 1: Interface and properties of version control system

Module:

Name: *GitClient* , **instance** *gc*.

Events:

Request: $\langle gc, \text{Push} \mid m, n \rangle$: Invokes a **push** operation of commit **m** with commit id **n**.

Request: $\langle gc, \text{Pull} \mid n \rangle$: Invokes a **pull** operation starting at commit id **n**.

Indication: $\langle gc, \text{PushReturn} \mid n \rangle$: Completes a **push** operation with commit **n**.

Indication: $\langle gc, \text{PullReturn} \mid m \rangle$: Completes a **pull** operation with commit list **m**.

Properties:

G1: Termination.

G2: Agreement. } “Total Order Broadcast” Module

Version Control System

Implementation

Algorithm 1:

Implements:

GitClient, **instance** *gc*.

Uses:

ReliableBroadcast, **instance** *rb*.
Consensus (multiple instances)

upon event $\langle gc, \text{Init} \rangle$ **do**
???

upon event $\langle gc, \text{Pull} \mid n \rangle$ **do**
???

upon event $\langle gc, \text{Push} \mid m, n \rangle$ **do**
???

upon event $\langle rb, \text{Deliver} \mid m, n \rangle$ **do**
???

upon event $undecided \neq \emptyset$ **and** $wait = \text{False}$ **do**
???

upon event $\langle c.r, \text{Decide} \mid \text{decided} \rangle$ **do**
???

Version Control System

Implementation

Algorithm 1:

Implements:

GitClient, **instance** *gc*.

Uses:

ReliableBroadcast, **instance** *rb*.
Consensus (multiple instances)

upon event $\langle gc, \text{Init} \rangle$ **do**

commits := *undecided* := \emptyset ;

round := 1

wait := **False**

upon event $\langle gc, \text{Pull} \mid n \rangle$ **do**

trigger $\langle gc, \text{PullReturn} \mid \text{commits}[n:\text{end}] \rangle$

upon event $\langle gc, \text{Push} \mid m, n \rangle$ **do**

trigger $\langle rb, \text{Broadcast} \mid \langle m, n \rangle \rangle$;

upon event $\langle rb, \text{Deliver} \mid m, n \rangle$ **do**

if $m \notin \text{undecided}$ **then**

undecided := *undecided* $\cup \{m, n\}$;

upon event *undecided* $\neq \emptyset$ **and** *wait* = **False** **do**

wait := **True**

trigger $\langle c.\text{round}, \text{Propose} \mid \text{undecided} \rangle$;

upon event $\langle c.r, \text{Decide} \mid \text{decided} \rangle$ **do**

forall $(m, n) \in \text{sort}(\text{decided})$ **do** // Sort by commit n° .

commits := *commits* $\cup m$; // Apply commit n° .

trigger $\langle gc, \text{PullReturn} \mid n \rangle$

undecided := *undecided* $\setminus \text{decided}$;

round := *round* + 1;

wait := **False**;

Content Sharing System

Problems:

- Which consensus algorithm would you use
 1. In a **fail-stop** system?
 2. In a **fail-silent** system?
 3. In a **byzantine** system?

- Is Uniform Consensus necessary? When ?

HOMework !