

Assignment 1: Vector Clocks

(1 P.)

- (a) Alice, Bob, and Carol are planning to have dinner on Saturday in a downtown restaurant. Alice proposes to check out the new Chinese restaurant and sends an email to Bob and Carol. Bob responds via Skype to Alice and Carol, proposing the Irish Pub instead; Alice accepts the suggestion and responds to Bob, but Carol did not receive Bob's message at all. So she (Carol) herself suggests to Alice and Bob that the Chinese restaurant is fine with her.
- (i) Describe the problem in a flow of communication processes between the three persons in a graphical way and specify the corresponding vector clocks at each change.
 - (ii) What is the final state of the "discussion". Is there a problem shown by the vector clocks? If, how can the problem be solved?
- (b) Properties of Vector Clocks (VCs) and relation to Lamport timestamps:
- (i) Formally show that $e_1 \rightarrow e_2 \Leftrightarrow VC(e_1) < VC(e_2)$ or give a counter example for both or one of the two directions.
 - (ii) Formally show that $VC(e_1) < VC(e_2) \Rightarrow C(e_1) < C(e_2)$, where $C(x)$ is the Lamport timestamp for event x .

Assignment 2: Consistency Models

(1 P.)

Consider a regular soccer gameplay of the home team against visitors. The goals are recorded in a NoSQL key/value store, one object for the home team (key: "home") and one for the visitors (key: "visitors"). When a team scores a goal the corresponding value is read from the key/value store, incremented by one and written back to the store. Consider the following writes to the key/value store and the minute at which the write was issued/goal happened.

Operation to key/value store	Minute of the Game
Write("home", 1)	12
Write("visitors", 1)	31
Write("home", 2)	37
Write("home", 3)	44
Write("visitors", 2)	56
Write("home", 4)	64
Write("home", 5)	68

- (a) The game is currently in the 69th minute. The key/value store runs on a cluster of several machines; data is replicated accordingly. Different consistency guarantees for the read operation can lead to different values, as discussed in the lecture. Fill in the possible reads in the table:

Consistency Model	Possible Values
Strong Consistency	
Consistent Prefix	
Bounded Staleness	
Eventual Consistency	
Monotonic Reads	
Read my Writes	

For the bounded staleness, consider a guarantee of 15 minutes.

(b) Consider the following three people

- The radio moderator that gives reports about the game every few minutes to his/her audience.
- The official scorekeeper at the game is responsible for tracking the goals scored. He is the one that updates (writes to) the key/value store.
- The sportswriter is going to write a newspaper article about the game.

Discuss the applicability of different consistency models for each of the tasks the individual people have to fulfill. Which is the appropriate model, which ones would lead to (what kind of) problems.

Assignment 3: HBase Characteristics

(1 P.)

Take a look at HBase and compare its wide-column store model to traditional relational databases and also to key-value stores like Riak. Describe one “ideal” scenario where it can be used and which benefits arise from its usage compared to the other two aforementioned possibilities.