

# Revision

# Key Words in Questions

Name	Usually a single word will do
Why?	1 or 2 sentences
Explain	1 or 2 sentences
Design	Long answer

# How does that impact marks?

Won't take marks off for longer answers.

... Brief will save you time for the design questions.

# How does that impact marks?

Won't take marks off for longer answers.

... Brief will save you time for the design questions.

... We do take off if you say something wrong.

# No multiple guessing

If the question says “name two” and you name three, we’ll mark the first two.

# Pick 2 Questions of 3

The exam has three questions: 1, 2, and 3 (each of which has parts).

Choose exactly two and indicate this on the front.

If you do all three questions, we will only mark 1 and 2.

... unless, in our judgment, one is obviously incomplete.

This is the short version of school policy.

# Difference from Last Year

Exams before academic year 2015–2016 are from a different instructor.

Less scatterbrain than last year.

# Emphasis on solving problems

Design a system to store YouTube videos.  
You're Twitter. Make a low-latency feed system.  
Show the most frequently visited pages. Live.

# Reason About a Problem

- What's the bottleneck? Compute, storage, bandwidth, latency, ...
- Which data is big? Which is small?
- Is it balanced?
- What if that machine fails?

# Reason About a Problem

- What's the bottleneck? Compute, storage, bandwidth, latency, ...
- Which data is big? Which is small?
- Is it balanced?
- What if that machine fails?

More fun: how will this system fail?

# Bag of Tricks

- Sharding aka partitioning: divide the work across machines
- Replication for speed and fault tolerance
- “Cold” large read-only store, “hot” small mutable store
- Approximations: Bloom filters, streaming counts, reservoir sampling

# Systems

**MapReduce** Parallel batch processing

**BitTorrent** File sharing

**HDFS** File storage

**Chord** Divide responsibility as machines join and leave

**SSTable** Large read-only key-value store

**BigTable** Large mutable key-value store

# Systems

MapReduce Parallel batch processing

BitTorrent File sharing

HDFS File storage

Chord Divide responsibility as machines join and leave

SSTable Large read-only key-value store

BigTable Large mutable key-value store

Given a problem, name a system and apply it.  
Take inspiration from the design of these systems.

# Web Service: Three Tiers

A loose design paradigm:

- 1 Stateless web servers
- 2 Application logic, high-level storage (i.e. BigTable)
- 3 Durable storage: Google File System, Lustre

# Starbucks Example

Starbucks lets gift card holders transfer funds between cards online.

- Tier 1** Route requests to Tier 2, render pages.
- Tier 2** Logic sharded by user. Ensure sufficient funds exist and decrement/increment balances of their users. Commits changes to Tier 3.
- Tier 3** Filesystem. Transfer records are appended to files and replicated across machines.

# State in Tier 1

Tier 1 machines are generally stateless.

Caching pages and even user information is OK.

Example: Google doodles on home page have a canned query, probably cached in Tier 1.

# Loosely Defined Tier 1 and 2

Starbucks credit moves from one user to another.

Tier 1 server could coordinate transaction with two Tier 2 machines.

Or Tier 2 server for sender contacts Tier 2 for receiver.

# Combine Tier 2 and 3?

Why doesn't the machine responsible for a user also store the final copy?

Fine for one application, but in general there are multiple applications sharing one filesystem.

# Example Exam Questions

## Setting the Scene

Starbucks issues reloadable gift cards. The website allows cardholders to transfer funds between cards. Cardholders have figured out that simultaneous transfers of money in and out of the same card sometimes results in more money in total.

# Example Exam Questions

## Setting the Scene

Starbucks issues reloadable gift cards. The website allows cardholders to transfer funds between cards. Cardholders have figured out that simultaneous transfers of money in and out of the same card sometimes results in more money in total.

## Applying Knowledge

a) Does the current Starbucks system guarantee linearisable consistency? [1 mark]

# Example Exam Questions

## Setting the Scene

Starbucks issues reloadable gift cards. The website allows cardholders to transfer funds between cards. Cardholders have figured out that simultaneous transfers of money in and out of the same card sometimes results in more money in total.

## Applying Knowledge

a) Does the current Starbucks system guarantee linearisable consistency? [1 mark]

No.

# Example Exam Questions

## Setting the Scene

Starbucks issues reloadable gift cards. The website allows cardholders to transfer funds between cards. Cardholders have figured out that simultaneous transfers of money in and out of the same card sometimes results in more money in total.

## Evaluating Technology

b) Would BigTable's consistency model solve the problem? Why?  
[2 marks]

# Example Exam Questions

## Setting the Scene

Starbucks issues reloadable gift cards. The website allows cardholders to transfer funds between cards. Cardholders have figured out that simultaneous transfers of money in and out of the same card sometimes results in more money in total.

## Evaluating Technology

b) Would BigTable's consistency model solve the problem? Why?  
[2 marks]

No, because it can't do transactions across rows.

# Example Exam Questions

## Setting the Scene

Starbucks issues reloadable gift cards. The website allows cardholders to transfer funds between cards. Cardholders have figured out that simultaneous transfers of money in and out of the same card sometimes results in more money in total.

## Design

c) Design a new replacement system to handle gift card balances and support funds transfers without creating additional money. Your system should follow the three-tier cloud model. For each of the three tiers, explain what the machines do and how their work is divided.[8 marks]

# Example Exam Questions

- c) Design a new replacement system to handle gift card balances and support funds transfers without creating additional money. Your system should follow the three-tier cloud model. For each of the three tiers, explain what the machines do and how their work is divided.[8 marks]
- Tier 1 is presentation with front-line web servers. These contact the Tier 2 machine for the user from whom funds will be transferred.

# Example Exam Questions

c) Design a new replacement system to handle gift card balances and support funds transfers without creating additional money. Your system should follow the three-tier cloud model. For each of the three tiers, explain what the machines do and how their work is divided.[8 marks]

- Tier 1 is presentation with front-line web servers. These contact the Tier 2 machine for the user from whom funds will be transferred.
- Tier 2 has machines responsible for shards of users, using Chord to assign responsibility. These machines ensure sufficient funds exist and decrement/increment balances of their users. The Tier 2 machine contacts the machine responsible for the destination to send funds. Agreement is done by two-phase commit (coordinated by Tier 1).

# Example Exam Questions

c) Design a new replacement system to handle gift card balances and support funds transfers without creating additional money. Your system should follow the three-tier cloud model. For each of the three tiers, explain what the machines do and how their work is divided.[8 marks]

- Tier 1 is presentation with front-line web servers. These contact the Tier 2 machine for the user from whom funds will be transferred.
- Tier 2 has machines responsible for shards of users, using Chord to assign responsibility. These machines ensure sufficient funds exist and decrement/increment balances of their users. The Tier 2 machine contacts the machine responsible for the destination to send funds. Agreement is done by two-phase commit (coordinated by Tier 1).
- Tier 3 is the backing store (filesystem) where Tier 2 ensures transaction records are committed.

# Linearisable versus Sequential

Alice and Bob write checks to each other for the same amount.

## Alice's Statement

-10 Check Alice → Bob  
0 Check Bob → Alice

## Bob's Statement

-10 Check Bob → Alice  
0 Check Alice → Bob

Both overdraft.

- ✓ Sequential: each client sees a consistent order
- ✗ Linearizable: no globally linear story