

EXAMINATIONS – 2013
TRIMESTER 2

<p>SWEN 433 Web Information Systems Engineering</p>
--

Time Allowed: THREE HOURS

Instructions:

- Closed Book.
- Answer all questions.
- Read each question carefully before attempting it.
- Total marks are 180.
- Each question is worth 60 marks, and should take about 60 minutes to answer.
- No calculators are permitted.
- Paper dictionaries for translating between English and a foreign language are permitted.

Question 1. WISE Foundations and Web Services

[60 marks]

(a) [15 marks] Suppose you are a software engineer working for Wellington's public transport provider *Metlink*. There are plans to build a new web information system for their business. You want to convince them that **web services** are the right technology to use.

- i. Summarize what the **core principles** of web services are.
- ii. Give **examples** for web services that could be developed and/or used at Metlink.
- iii. Briefly explain what **service descriptions** are, and why they are needed.

(b) [15 marks] Discuss what the **REST architectural style** is. *Elaborate on the role of resources in REST.* What are the benefits of REST?

(c) [15 marks] Compare and contrast **big web services** to **RESTful web services**. *Discuss advantages and disadvantages, and give examples of scenarios when you would favour one over the other.*

(d) [15 marks] Explain what **quality of service** is, why it is important for web services, and how it is related to **service-level agreements**.

Question 2. Web Data Management

[60 marks]

(a) [15 marks] Compare **horizontal** and **vertical fragmentation**. *Why do the disjointness rules for these two fragmentation approaches differ?*

(b) [15 marks] Discuss the pros and cons of using **data replication** in a web information system. *What additional tasks must be taken care of by a web information system for replication management?*

(c) [15 marks] Explain what the **CAP theorem** says. *Discuss the major consequences for web information systems.*

(d) [15 marks] Consider the following scenario: The server at a travel agency initiated a distributed transaction involving four other **participants**: an international airline (P_1), a hotel chain (P_2), a car rental agency (P_3), and an insurance company (P_4). The travel agency acts as the **coordinator** (C) in this transaction and can communicate directly with all other participants. Assume that all parties agree to the transaction and are willing to commit.

- i. Assuming that there are no failures, describe the sequence of messages exchanged and the log entries written at each node when using the **2PC protocol** to perform this transaction.
- ii. Now, suppose that the server of participant P_2 crashes just after sending the vote-commit message, and there are no further failures. Hence, the coordinator will still successfully receive this message. For this modified situation, answer the question from Part i. Will any of the participants be blocked?

Question 3. NoSQL Databases

[60 marks]

(a) [15 marks] Do you think that **NoSQL databases** will replace **relational databases** in web information systems? *Justify your answer.*

(b) [15 marks] Many NoSQL databases use BASE for **transaction processing**. Explain what **BASE** is, and how it is different from the classical ACID properties of transactions. *Give examples to illustrate when BASE is a good choice.*

(c) [15 marks] Compare **document stores** and **key-value stores**. *Discuss similarities and differences.* When would you recommend the use of these two categories of NoSQL databases? *Give examples.*

(d) [15 marks] Many NoSQL databases like BigTable utilize a **shared-nothing architecture**. *Explain how this is different from a shared-disk architecture, and how this affects scalability.*
