

Advanced Database Design and Implementation - Course Outline

SWEN 432: 2014 Trimester 1

This document sets out the workload and assessment requirements for SWEN 432. It also provides contact information for staff involved in the course. If the contents of this document are altered during the course, you will be advised of the change by an announcement in lectures and/or on the course web site. A printed copy of this document is held in the School Office.

SWEN 432 is an advanced course in database and information systems. Databases and information systems facilitate the handling of large amounts of information. The information we want to process is now more complicated than it used to be. We require more intelligent, application-specific, and sophisticated systems to do the processing.

This year the course will examine the following two contemporary fields in the database systems area:

- XML Data Model and XML Databases, and
- Cloud Databases.

XML Data Model and XML Databases will comprise approximately 70% of the course. There, we shall consider topics such as: XML documents, Document Type Definition (DTD) and XML Schema, XML constraints, XML query languages, Types of XML Databases, Mapping XML data to relational databases, Publishing relational databases as XML documents, and what research is going on in the XML database area. The practical experience will be achieved through the use of XML processors like `xmllint` and the native XML database management system `eXist`.

The Cloud Database field will comprise approximately 30% of the course and will consider questions like: what is Database as a Service (DaaS), what are main features of cloud databases and how they differ from traditional databases, and what are CAP and BASE trade-offs in cloud database implementations. Further, three specific data models used to deploy cloud database management systems will be considered. There will be also presented details of three cloud database implementations (Amazon's Dynamo, MongoDB, and Apache's Cassandra). A practical assignment will use Cassandra to demonstrate issues in designing, building and querying cloud databases relying on the column data model.

Objectives

By the end of the course, students should be able to:

1. Design well formed XML documents that are valid with regard to a given DTD or XML Schema and thus develop the ability to solve practical engineering problems ([BE graduate attribute 3\(f\)](#)),
2. Analyze a part of the real world and design a corresponding XML DTD or Schema in XML normal form and thus develop the ability to formulate and build efficient models of complex systems using principles of engineering science and mathematics ([BE graduate attribute 3\(b\)](#) and [BE graduate attribute 3\(c\)](#)),
3. Design faithful models of a part of the real world using XML database constraints and thus develop the ability to apply mathematical and engineering science in solving engineering problems ([BE graduate attribute 3\(a\)](#)),
4. Use available web sources to learn about the `eXist` XML database management system and define XQuery queries and XUPdata updates of an almost arbitrary complexity against a native XML database and thus develop the ability to look for additional information from pertinent sources ([BE graduate attribute 3\(d\)](#)),
5. Analyze properties and usage of a set of XML documents, decide on the most appropriate kind of a database management system and implement the database and thus develop ability to formulate the problem and create a solution that models behavior of a part of the real world ([BE graduate attribute 3\(c\)](#)),
6. Demonstrate a comprehensive understanding of Cloud Database structures and mechanisms they use to achieve high availability, scalability, network partition tolerance, and eventual consistency.
7. Analyze a set of user queries and design a column oriented cloud database using CQL, the data language of Cassandra,
8. Implement a cloud database and query it and thus develop the ability to synthesize a solution and make design documentation ([BE graduate attribute 3\(b\)](#) and [BE graduate attribute 2\(b\)](#)), and
9. Read, critically analyze, and write reports on research articles in an area covered by the course and thus develop the ability to communicate effectively ([BE graduate attribute 2\(b\)](#)).

Knowledge assumed of participants

Before you enroll in SWEN 432 you are expected to have passed SWEN304 OR COMP302. Particularly, you will need to be familiar with the following topics:

- Relational Data Model,
- Structured Query Language (SQL),
- Relational Functional Dependencies and Normal Forms,
- PostgreSQL Data Base Management System.

You must also have a good understanding of English and be able to write clearly.

Recommended Reading

Have a look at [reading links](#). It is a reading list of material that may be of use. You will also find there a link to PostgreSQL and eXist documentation.

Lectures

A [schedule](#) of lecture topics and tentative assignment due dates is available online. Lectures for SWEN 432 are:

- Tuesday 12:00 to 12:50 am in MY403,
- Wednesday 12:00 to 12:50 am, in MY403,
- Friday 12:00 to 12:50 am, in MY403.

Assignments

There will be five assignments. A tentative assignment schedule is given [here](#). In principle, each assignment will be issued on Monday and due on Friday next week. In your answers to assignment questions, you are expected to communicate your solutions in a clear way and, if needed, to present a proper design documentation. If you think you need more information than covered in lectures, use [reading links](#), where you will find XML standards, and links to other sources.

The first assignment covers basics of XML documents and their two meta languages DTD and XML Schema, helping you to meet objectives one and two. The second assignment is about XML Schema identity constraints, XML functional dependencies and XML normal form. The second assignment brings into focus aspects of designing high quality XML schemes using mathematically grounded methods and thus to achieve the course objective three. The third assignment considers querying and updating XML documents using the native XML database management system eXist. The third assignment helps you to achieve objective four. In the fourth assignment you are asked to design and implement the same XML database using three different approaches and to compare their performance. The fourth assignment leads to the course objective five. Finally, the assignment five asks you to design, and build a small cloud database and to exercise querying it using Cassandra. By doing the assignment five, you achieve the course objectives six, seven, and eight.

Assignments will be handed out in lectures. You will find all necessary information about assignments and files needed [here](#). Assignments should be submitted [electronically](#). Any assignment submitted after the due date will be penalized at the rate of 5% per day (a weekend counts as one day). Approval to submit assignments late without penalty will only be granted in exceptional circumstances and if the arrangement is made prior to the due date. Any medical excuse must be accompanied by a doctor's certificate.

Essay

Writing an essay is an important part of your assessment. The essay helps you to achieve the course objective nine. You will find the essay title and description in [essay 2014 guidelines](#). The essay is *due on Friday 23 May*. You may find it useful to consult also [general](#) essay guidelines. Your essay should not exceed 7000 words. Submit the essay electronically as a .ps, or .pdf file, and also print it and hand in a hard copy into the hand-in box. There will be no essay presentations this year. The mark you get for your essay will be based on how interesting the content of your essay is, how readable the essay is, and how good your analysis of the problem is.

Workload

In order to maintain satisfactory progress in SWEN 432, you should plan to spend an average of at least 10 to 12 hours per week, which includes attending lectures, solving assignments, writing essay, reading literature and reviewing lecture material.

School of Engineering and Computer Science

The School office is located on level three of the Cotton Building ([Cotton 358](#)).

Staff

The course organizer and lecturer for SWEN 432 is: Pavle Mogin. His contact details are:

- [Pavle Mogin](#)
- Room: [Cotton 233](#)
- Phone: 463 5443
- email: Pavle.Mogin@ecs.vuw.ac.nz
- Office hours: Tuesday 2 to 3 pm

[Aaron Morton](#) is going to deliver three lectures on Cassandra; a cloud database management system, as a guest lecturer.

Your class representative is:

- Person: Pragya Mohan

- email: mohanprag@gmail.com

Announcements and Communication

The main means of communication outside of lecture will be email and the SWEN 432 web area at http://ecs.victoria.ac.nz/Courses/SWEN432_2014T1/. There you will find, among other things, this document, the [lecture schedule](#), [assignment handouts](#), and the [SWEN 432 Forum](#). The forum is a web-based bulletin board system. Questions and comments can be posted to the forum, and staff will read these posts and frequently respond to them.

Assessment

Your grade for SWEN 432 will be determined from your work on the essay, assignments, and the final examination, as follows:

Item	Weight	Week Due
Essay	12%	10
Assignment 1	4%	4
Assignment 2	5%	6
Assignment 3	5%	7
Assignment 4	4%	9
Assignment 5	5%	12
Final examination	65%	

Each assignment contains a number of questions worth 100 marks in total. Your answers to questions will be assessed according to their accuracy. You may find more about the assessment of the essay in [essay 2014 guidelines](#).

Final Examination

The [timetable for final examinations](#) will be available from the University web site and will be posted on a notice board outside the faculty office. The final examination will be three hours long. No computers, or similar device will be allowed in the final examination. Only paper non-English to English dictionaries and simple calculators will be permitted. The study and examination period for trimester T1 is 9 June to 02 July.

Plagiarism

Working Together and Plagiarism We encourage you to discuss the principles of the course and assignments with other students, to help and seek help with programming details, problems involving the lab machines. However, any work you hand in must be your own work. The [School policy on Plagiarism](#) (claiming other people's work as your own) is available from the course home page. Please read it. We will penalize anyone we find plagiarizing, whether from students currently doing the course, or from other sources. Students who knowingly allow other students to copy their work may also be penalized. If you have had help from someone else (other than a tutor), it is always safe to state the help that you got. For example, if you had help from someone else in writing a component of your code, it is not plagiarism as long as you state (e.g., as a comment in the code) who helped you in writing the method.

Mandatory Requirements

To pass SWEN 432 a student must:

- hand in an essay, and
- achieve at least 40% for the assignments.

Passing SWEN 432

To pass SWEN 432, a student must satisfy mandatory requirements, achieve at least a D grade for the exam, and gain at least a C- grade overall.

Course Withdrawal

The last date for withdrawal from SWEN 432 with entitlement to a refund of tuition fees is Friday 14 March 2014. The last date for withdrawal without being regarded as having failed the course is Friday 16 May 2014 -- though later withdrawals may be approved by the Dean in special circumstances.

You will no longer be liable for the applicable course fees if you apply for a withdrawal by completing a Change of Course form (available from your Faculty Student Administration Office) or apply in writing on or before Friday, 15 March 2013.

Rules & Policies

Find key dates, explanations of grades and other useful information at <http://www.victoria.ac.nz/home/study>.

Find out about academic progress and restricted enrolment at <http://www.victoria.ac.nz/home/study/academic-progress>.

The University's statutes and policies are available at <http://www.victoria.ac.nz/home/about/policy>, except qualification statutes, which are available via the Calendar webpage at <http://www.victoria.ac.nz/home/study/calendar> (See Section C).

Further information about the University's academic processes can be found on the website of the Assistant Vice-Chancellor (Academic) at <http://www.victoria.ac.nz/home/about/avcacademic>

All students are expected to be familiar with the following regulations and policies, which are available from the school web site:

[Grievances](#)

[Student and Staff Conduct](#)

[Meeting the Needs of Students with Disabilities](#)

[Student Support](#)

[Academic Integrity and Plagiarism](#)

[Dates and Deadlines including Withdrawal dates](#)

[School Laboratory Hours and Rules](#)

[Printing Allocations](#)

[Expectations of Students in ECS courses](#)

The School of Engineering and Computer Science strives to anticipate all problems associated with its courses, laboratories and equipment. We hope you will find that your courses meet your expectations of a quality learning experience.

If you think we have overlooked something or would like to make a suggestion feel free to talk to your course organiser or lecturer.

[Course Outline as PDF](#)
