

## Course Change Proposal for Advanced Databases (INFR11011)

I would like to make small adjustments in the course description for Advanced Databases (INFR11011).

- 1) The course content remains mostly unchanged. I have updated the course description to remove references to the things I do not teach (e.g., indexing of multi-dimensional data) and add few new topics (e.g., multi-version concurrency control).
- 2) The assessment format now consists of two summative programming assessments and (at least) two formative assessments.
- 3) I have also updated the learning outcomes and other requirements to make them more precise and in line with School policies and practice.

Below you can find the DRPS sections affected by these changes (the other sections remain unchanged).

Best regards,  
Milos

### **Postgraduate Course: Advanced Databases (Level 11) (INFR11011)**

#### **[Summary]**

This course studies the data structures and algorithms underlying modern database management systems. It covers a wide range of topics and techniques such as storage architecture, data models (row, columnar), indexing (tree-based, hash tables), transaction processing (ACID, concurrency control), crash recovery, and parallel architectures (multi-core, distributed).

#### **[Course Description]**

This course will cover the following topics:

- Data models and data management systems architectures
- Storage: secondary-storage devices
- Indexing: tree-based and hash-based techniques
- Query evaluation: sorting and join processing, selection, projection, aggregation
- Query optimisation: cardinality estimation, cost-based query optimisation, dynamic programming, rule-based optimisation
- Transaction management: ACID properties, concurrency control, locking and multi-version protocols, crash recovery
- Distributed database systems: parallel query evaluation, distributed transaction processing.

**[Other Requirements]**

This course is open to all Informatics students including those on joint degrees. For external students where this course is not listed in your DPT, please seek special permission from the course organiser.

The course assumes an understanding of algorithms and data structures (e.g., quick sort, merge sort, binary trees, hash tables, big-O notation).

A good level of programming is assumed and will not be covered during lectures. The coursework will involve implementing query processing algorithms in the context of a database system written in C++.

**[Assessment]**

Written Exam 70%, Coursework 30%, Practical Exam 0%

**[Additional Information (Assessment)]**

Two programming assignments, each worth 15%.

You should expect to spend approximately 50 hours on the coursework for this course.

**[Feedback]**

Formative feedback: At least two sets of homework assignments submitted during the semester with feedback returned within two weeks, and feedback on exam papers.

**[Learning Outcomes]**

On completion of this course, the student will be able to:

- Understand how database management systems function internally.
- Interpret and comparatively criticise database systems architectures.
- Implement major components of a database management system and analyse their performance.
- Understand, analyse, and compare the fundamental query evaluation and concurrency control algorithms.
- Identify trade-offs among database systems techniques and contrast distributed/parallel techniques for OLTP and OLAP workloads.

Milos Nikolic

Lecturer in Database Systems

School of Informatics, University of Edinburgh

<http://homepages.inf.ed.ac.uk/mnikolic/>