

EECS 484 - Fall 2020

Database Management Systems

Instructors:

Kevin Leach and Barzan Mozafari

Course Staff Email: 484questions@umich.edu

GSI: Sahil Farishta;

IAs: Nick Holden, Junhwan Kim, Gabriel Garfinkel, Kevin Sang, Matthew Zhang, Yuning Wang

Course Description

Databases are extremely prevalent -- they are behind technology that all of us use daily, such as mobile apps, banking systems, video games, telecommunication systems, course registration, supermarket transactions, movie and music services, and many more. This course gives a basic introduction to database management systems (DBMSs), i.e. software that is designed to assist in maintaining and utilizing large collections of data. The focus will be on relational DBMSs, which are by far the dominant type of DBMS. The course is designed to provide you with both an external and an internal view of relational DBMSs. By the end of this course, you will know:

- How to design a DBMS, store data for real-world scenarios, and develop applications that communicate with the DBMS.
- How to analyze data by "asking" appropriate queries.
- The internal design of DBMSs that makes them robust to system failures and supports concurrent data accesses.
- How DBMSs are able to store massive datasets and support efficient query handling.

Prerequisites

Prerequisites EECS 281 or graduate standing: You are required to have the background from a data structures course (e.g., lists, hash tables, arrays, search trees) and strong programming experience.

Programming skills: The projects are in Java and C++. The knowledge of C++ is assumed. You will be given a brief overview of Java and expected to be able to pick up the language on your own, if you haven't used it before.

Office Hours

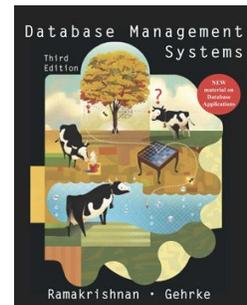
The **office hours** of the teaching staff will be available on the [course calendar](#) along with the room number or the Zoom details. Before heading to the office hours, always make sure to check the latest schedule on the calendar or piazza posts for exceptions.

Course Materials

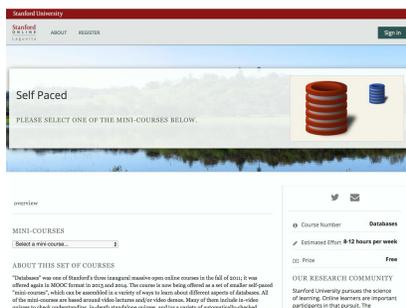
- **Piazza:** We will use Piazza for Q&A. In addition, announcements, lecture recordings, slides, etc. will all be posted via Piazza. Please monitor the course Piazza regularly: <https://piazza.com/class/ke05dmqen9i1lr>
- **Autograder.io:** Your projects will be graded via the EECS autograder.io service: <https://autograder.io>
- **GradeScope:** You will submit written homework assignments via GradeScope: <https://gradescope.com>
- **Canvas:** We will *not* be using Canvas this semester.

Textbook & Other Resources

We will follow the “Cow” book, i.e., “Database Management Systems” (3rd edition), by Raghu Ramakrishnan and Johannes Gehrke, McGraw Hill, 2003, ISBN: 0-07-115110-9. The book has a supporting website (<http://pages.cs.wisc.edu/~dbbook/>), where you can find answers to odd-numbered exercises. **The book is really useful**, especially for exams. **We follow the book pretty closely.**



In addition, **we recommend** the following informal reference: “**Self-paced courses on introduction to Databases**” by Jennifer Widom, Stanford University



<https://www.edx.org/course/databases-5-sql>

There are several mini-courses at the above link and they are currently free. They are all very useful from a practical perspective and nicely complement the material that we cover in EECS 484.

Course Announcements

We will use **Piazza** for course announcements. You are responsible for ensuring that you receive notifications from Piazza and that you pay attention to any announcements. The default settings will send you email notifications every 2 hours -- we recommend configuring Piazza to deliver at least daily digests so that you do not miss useful information.

You should be enrolled in the Piazza: <https://piazza.com/class/ke05dmqen9i1lr>. If you do not have access, please email the instructional staff to be added.

Posting Questions

You must **strictly** follow the rules below regarding email communications with the instructors and other staff members of the class, or else your email will be ignored:

- Questions about **homeworks and projects, and technical questions: Use Piazza**.
 - You should first consider posting your question there, AFTER **searching to see if a similar question has already been answered**.
 - **You should strictly refrain from sending private questions UNLESS** posting your question is going to be a violation of honor code (e.g., you are posting something from your solution). With non-private questions, other students will also benefit from your answers.
 - **Giving answers:** We encourage you to give answers to conceptual questions by other students. The teaching staff may "endorse" your answer or give their own answer. Of course, be careful of the Honor code---you don't want to solve a homework problem for someone, only help them conceptually so that they are on the right track.
 - **Be contributive and helpful:** If you post a question and figure out the solution on your own, please share your solution to the whole class by writing it down in student's answer or follow-up discussion instead of delete the post or put "resolved" in the answer.
- **Personal requests and special accommodations:** For questions that are definitely not relevant to other students, send email to 484questions@umich.edu. This will reach all the teaching staff. If your question is relevant to a specific TA (not the instructor), feel free to email the appropriate address, but make sure to **prefix the subject of email with [EECS 484]**. However, emailing 484questions@umich.edu is preferred, as then anyone can answer it and it is less likely to be overlooked (some of us do get a lot of emails!). We may not respond if you email personal addresses instead.

Given the large number of students in the class, by following the communication guideline, you will help us better serve your needs throughout the semester.

Oracle and Postgres Accounts

The programming assignments require you to use Postgres or Oracle accounts. We will provide instructions on obtaining the accounts in the discussions.

Course Requirements and Grading

Midterm Exam	20%
Final Exam	20%
Projects: 4 programming assignments, 12.5% each	50%
Homeworks: 5 written assignments, 2% each	10%
Bonus/Negative points (see below)	+/-2%

The examinations are ***not*** cumulative. Each covers one half of the course. In particular, the final exam covers material presented after the topics for the midterm (with the **exception of Tree-based Hashing** which will be tested in both exams).

To get a C or higher, passing performance is expected on both exams and projects. That is, you must earn a passing average on the exams in addition to a passing average on the projects. By default, “passing performance” means 70%+, though we may adjust this based on class average.

The midterm exam is scheduled for **Monday, October 19**. You will have the opportunity to take the exam during any two-hour window between the hours of 8AM-11:59AM, 2PM-5:59PM, or 8PM-11:59PM (all Eastern time) on that day. We will schedule IAs to be available during these times for clarification questions.

The final exam is scheduled for **Friday, December 11**. You will have the opportunity to take the exam during any two-hour window between the hours of 8AM-11:59AM, 2PM-5:59PM, or 8PM-11:59PM (all Eastern time) on that day. We will schedule IAs to be available during these times for clarification questions.

For either exam, you must let us know by **September 30** if you have SSD forms and require extended exam times. In such cases, you will have a 3- or 4-hour window (as appropriate) within the 3 times listed above for each exam day.

We will not offer make-up exams. Make sure that you do not have any conflicts on the days of the midterm and final. Exams will be delivered in blocks of time according to the course schedule. Only in documented exceptional circumstances (subject to instructor pre-approval) will we consider alternate exam times.

Bonus/Negative points: Constructive contributions to answering other students' questions on Piazza and class participation can be rewarded by up to +2% as bonus points (at the instructor's discretion). Likewise, disruptive behavior such as rude or disruptive behavior on Piazza will be penalized by up to -2% (instructor discretion).

Assignments and Partners

We *strongly recommend* that **projects** be completed in **groups of two**. If you cannot find a group, the instructional staff will assign you one. Only under special circumstances, and with explicit permissions granted by the teaching staff, will working alone be permitted. If you decide to work alone on the projects, please be aware that your workload will be higher. All **homeworks** must be done individually.

File types for submission:

Normally, most submissions will require you to either submit one or more source code or SQL files (for projects) to the autograder or a single PDF (for homeworks) to GradeScope. The autograder will tell you what file(s) you are supposed to upload for each project.

Project Submissions

This course will use the EECS autograder.io service (<https://autograder.io/web/course/77>). For each project, you will register as a group, and get 3 submissions per day. We will accept the submission that receives the highest grade. It is against the Honor Code to register individually on the autograder while working in a group. We will use automated plagiarism detection software on your submissions. Aside from high-level discussion, collaborating with others outside of your group is not permitted.

Your submissions should indicate the contributions made by each individual in your submission. We may make adjustments to your individual grade after project submissions based on individual contributions.

Homework Submissions

We will use GradeScope for HW submissions: <https://www.gradescope.com/courses/164130>. You should have been added to the GradeScope roster. Email course staff if you are not enrolled.

Late Days

For the homework and project submissions, check out the schedule on the website. Typically, they will be due on **Wednesdays at 11:55pm**.

For **homeworks**, you will get a **zero on missing the deadline**. We recommend submitting a version several hours before and verifying that the autograder.io service actually shows the correct files as submitted (**it is your responsibility to verify submissions**).

For each project, there a 4-day late submission period (i.e., if the project is due on Wednesday, 11:55 PM, a late submission is due before Sunday, 11:55 PM) in which a **15% late penalty will be incurred for that project** (i.e., 1.5/10 points will be deducted, minimum score of 0). Beyond that, you will get a zero on the project. Please submit at least 30 minutes before the regular deadline as a safety measure. **Late days are rounded up to the nearest integer**. For example, a submission that is 4 hours late will count as one day late.

We have run into situations in the past (rare) where a student misses the regular deadline by 2-3 minutes for a project and incurred a 15% penalty. Sometimes, this is because of last-minute project work or slow servers. **We will give a one-time waiver of the 15% penalty if you miss the regular submission deadline for a project by 5 minutes or less** (i.e., 12:00 AM or earlier). **Beyond that, the 15% penalty will apply, even if you miss the deadline by 1 minute.** **We will not make ANY exceptions to our late submission policy.**

Honor Code

All students (including LS&A and Engineering) are required to observe the Engineering Honor

Code in all assignments and exams. A copy of the honor code can be found at <http://ossa.engin.umich.edu/honor-council/>. Please make sure that you clearly understand what constitutes cheating. If you are not sure in any specific case, you should ask the teaching staff. The University takes honor code violations seriously, and penalties can be severe. You are not allowed to share your code with anyone other than your partner. You are not allowed to make use of project or homework solutions by others, including solutions from previous semesters. **Make sure that you do not upload your code on public repositories (e.g., GitHub), as this also constitutes violation of the honor code.**

Any suspected violations of the honor code will be reported.

Disabilities and Conflicts

Students with disabilities that are documented with the Services for Students with Disabilities (SSWD) Office should contact the course staff during the first three weeks of class to make appropriate arrangements.