

University at Buffalo

CSE462/562 Database Systems (Fall 2024)

Instructor: Zhuoyue Zhao, zzhao35@buffalo.edu.

Lecture location & time: Tuesday and Thursday 5:00 pm to 6:20 pm, Nsc 205.

Instruction Mode: Lectures.

Deliver Mode: In-person.

TA: Congying Wang.

Grader: Gianna Bossoreale.

Office hours:

- Tuesday, 11:00 AM - 1:00 PM, Davis 338Y (Congying)
- Wednesday, 1:00 PM - 5:00 PM, Davis 338I (Zhuoyue)
- Thursday, 10:00 AM - 12:00 PM, Davis 338Y (Congying)
- No office hours on the following days: 8/27, 8/28, 8/29, 10/15, 11/27, 11/28
- Office hours ends prior to the final exam week.
- Subject to change. Please see Piazza and course website for updates.

Course website: https://cse.buffalo.edu/~zzhao35/teaching/cse562_fall124/

Announcement and discussion board: We will be using Piazza for announcements, discussion and Q&A outside the lecture and office hours: <https://piazza.com/buffalo/fall124/cse462562/home> **Number of**

Credits: 3

Course Description: Database Management Systems (DBMS) are an important piece of software widely used in various data intensive applications. No matter what data model or query language it adopts, there are principles and methodologies commonly found in virtually all DBMS, in order to support efficient and fault-tolerant data storage, query and management. In this course, we will examine the internals of a traditional Relational DataBase Management System (RDBMS) and discuss the common principles and methodologies that may be useful in many other database and data processing systems beyond RDBMS. The students will also gain practical experience of efficient data management through a semester-long project of implementing various components of a mini RDBMS.

The main topics include database storage management, query processing, query optimization, transaction processing, concurrency control and recovery in RDBMS. Please note that this is not a course on database design or database application development – those are covered in CSE460/560 Data Models and Query Languages. We will briefly review the necessary background on the relational data model and the SQL language in the first a few lectures but students are expected to understand relational algebra and SQL before this course.

Learning Outcomes:

The students should demonstrate mastery of the concepts/skills/knowledge expressed in the following ABET learning outcomes for computer science:

- (2) Design, implement and evaluate a computing-based solution to meet a given set of computing requirements in the context of the program's discipline.
- (5) Function effectively as a member or a leader of a team engaged in activities appropriate to the program's discipline.
- (6) Apply computer science theory and software development fundamentals to produce computing-based solutions.

Upon completion of this course, a student is expected to

- Be able to design and implement Relational DataBase Management Systems with basic storage and query processing capabilities (ABET 2; **Instructional Method:** Lectures; **Assessment Methods:** Homeworks, Exams, Projects);
- Be able to design and analyze efficient algorithms and data structures for data management systems (ABET 6; **Instructional Method:** Lectures; **Assessment Methods:** Homeworks, Exams, Projects);
- Be able to work in a team effectively for a complex software system project (ABET 5; **Instructional Method:** Lectures; **Assessment Methods:** Projects).

The Student Outcomes from the Computing Accreditation Commission (CAC) of the ABET have been adopted.

Program Outcome Support (Computer Science ABET Outcomes):

Program Outcome	1	2	3	4	5	6
Support Level	No coverage	Demonstrate mastery of skill/concept	No coverage	No coverage	Demonstrate mastery of skill/concept	Demonstrate mastery of skill/concept

Course Prerequisites:

Required: Solid background in programming and data structures. Programming experience with C++. For CSE462 students, CSE460 is a mandatory prerequisite.

No mandatory textbook: We will make all the lecture slides available on the course website.

Optional textbook:

- Database System Concepts, Seventh Edition, by Abraham Silberschatz and Henry Korth and S. Sudarshan.

Course Requirements:

- You are expected to understand and follow the academic integrity policy of UB, CSE and this course. We will *strictly* enforce the academic integrity policy (see below).
- In-person attendance is *required*. You should attend all the lectures and be familiar with all the course materials presented in class, which may appear in any course projects and exams. We will try to make the recording of the lectures available, but we cannot guarantee the availability of those in case of technical difficulties. Please note that there is no live streaming of the lectures.
- There will be a mid-term exam and a final exam. They must be attended in person and there will be no make-up exam except for the exam/course conflicts made known to the instructor within two weeks of the fall semester (9/9/2024), or unusual circumstances on a case-by-case basis. The exams are open-book but you may only bring the paper-copy of lecture slides, the written assignments, the textbook or your lecture notes. You should not use any electronic device during the exam, except for a calculator that may be needed.
- There is a semester-long course project of building a mini RDBMS in C++, divided into six subprojects. You may work individually or collaborate in a team of two. The tasks in each project will include a coding component, which you may collaborate with your teammate, and a write-up, which you must complete independently without any collaboration or sharing with your teammate. The tasks of the projects remain the same regardless of the team size, and you will receive the same grade for the coding component of your team. The write-up will be graded individually. We will use Autolab for code submission and UBLearn for write-up submission. Please refer to the course website for details.
- There will be four graded written assignments. The written assignments must be completed independently. Collaboration with classmates (including your project teammate) or using any materials other than the lecture slides and the optional textbook is not allowed.
- Late submission policy: you have up to 3 grace days in total for projects and written assignments, and you may only use at most 1 grace day for each project or written assignment. There is no penalty for late submissions that fall in the allowed grace days. No credit will be given if you make a

late submission that uses more than allowed grace days. If you have a project teammate and any of you make a late submission, any grace days used will count towards the used grace days of both team members. Both team members must have remaining grace days to count an allowed late submission towards your final grade.

Important Dates:

- Add/Drop deadline: September 3, 2024.
- Mid-term exam: October 17, 2024, 5:05 pm - 6:20 pm, Nsc 205.
- Last day to resign from the course: November 13, 2024.
- Final exam: December 12, 2024, 7:30 pm - 9:00 pm, Nsc 109.
- Project due dates: please refer to the course website.

Grading Policy:

The final grade will be broken down as follows:

- Mid-term exam: 15%.
- Final exam: 20%.
- Written assignments: 20%, 5% each.
- Projects: 45%.

And the following is the assignment of letter grades:

- $[90, +\infty)$: A
- $[80, 90)$: A-
- $[70, 80)$: B+
- $[60, 70)$: B
- $[50, 60)$: B-
- $[40, 50)$: C+
- $[30, 40)$: C
- $[20, 30)$: C-
- $[10, 20)$: D
- $[0, 10)$: F

The grades will not be curved and will not be differentiated between the sections.

Course Schedule:

The following is the course schedule may be subject to adjustment throughout the semester. Please refer to the course website for the latest schedule.

Lectures

Week#	Topics
1	Course introduction and logistics Relational model and query language
2	Physical storage, POSIX I/O interface, buffer management Data storage layout
3	Relational Algebra and SQL
4	Project and homework review Query processing overview
5	Volcano execution model and cost analysis Single-table query processing: selection, projection and expression evaluation
6	Single-table query processing: aggregation, sorting and hashing
7	Single-table query processing: deduplication and grouping Project and homework review
8	Midterm exam
9	Access methods and indexing Tree indexes
10	Tree indexes (cont'd) Hash Index
11	Index scans Concurrency control and recovery: an overview of transaction
12	Project and homework review Join algorithms
13	Join algorithms (cont'd) Query optimization
14	Query optimization (cont'd)
15	Project and homework review Advanced topics (tentative: random sampling and approximation in databases)
16	Final exam

Written Assignments

Written Assignment	Release Date	Due Date
HW1: relational model and query language	8/27/2024	9/16/2024
HW2: single-table query processing	9/19/2024	10/9/2024
HW3: access method, indexing and transaction	10/22/2024	11/11/2024
HW4: join processing and query optimization	11/14/2024	12/2/2024

Projects

Assignment	Release Date	Due Date
Project 0: Project and team sign-up	1/29/2024	2/1/2024
Project 1: POSIX I/O	9/3/2024	9/9/2024
Project 2: Data Layout	9/5/2024	9/16/2024
Project 3: Single-table query processing	9/19/2024	10/9/2024
Project 4: B-tree index	10/22/2024	11/11/2024
Project 5: Join algorithms and optimization	11/14/2024	12/2/2024

Academic Integrity Policy:

Academic integrity is critical to the learning process. It is your responsibility as a student to complete your work in an honest fashion, upholding the expectations your individual instructors have for you in this regard. The ultimate goal is to ensure that you learn the content in your courses in accordance with UB's academic integrity principles, regardless of whether instruction is in-person or remote. Thank you for upholding your

own personal integrity and ensuring UB's tradition of academic excellence. You should get familiar with the departmental and the university academic integrity policies and procedures:

- Undergraduate Academic Integrity Policy
- Undergraduate Academic Integrity Procedure
- Graduate Academic Integrity Policy
- Graduate Academic Integrity Procedure

Academic integrity policies specifics to this course (please read carefully): you may NOT share code with from anyone about your course projects except for your teammate. We also require all students, whether enrolled, dropped or resigned from the class, to keep your course project repository inaccessible to public indefinitely, and never share it with any current students or future students who may take the course. When you look up online materials for the course project, please exercise common sense in terms of what is allowed and what is not allowed (e.g., cplusplus.com are fine to look at, while posting questions directly related to the project on StackOverFlow would not be acceptable). Examples that we consider as academic integrity violations include but are not limited to:

Copying any part of other team's code implementation or code found online in the course project, regardless of whether it is a verbatim copy or a modified copy; Collaborate with other teams in the course project; (Writing code together with students from other teams or exchanging code is not allowed. Discussion on task and system specification, algorithms and system designs, on the other hand, is allowed and encouraged); Making your course project repository or any copy of it available to any current or future students of the course; Submitting work that is not created by you or your team, including deriving your solutions based on materials from search engines, tutoring services, Q&A websites, and/or those generated by generative AI software (e.g., ChatGPT); Collaborate with anyone in the written assignments; (Developing solutions together with fellow classmates or exchanging solutions is not allowed. Discussion on homework problems, course materials, on the other hand, is allowed and encouraged). Cheating or referring to any material not permitted in the exams and written assignments.

Consistent with the CSE Academic Integrity Policy, we have zero tolerance towards academic integrity violations. Any academic integrity violation will result in an F grade for all students involved and will be referred to the Office of Academic Integrity. The only exception when the violation is accidental and does not provide any unfair advantage to any of the students involved, in which case you may receive a reduction in grade in the particular assignment/exam/project.

Campus resources:

Accessibility Resources. If you have any disability which requires reasonable accommodations to enable you to participate in this course, please contact the Office of Accessibility Resources in 60 Capen Hall, 716-645-2608 and also the instructor of this course during the first week of class. The office will provide you with information and review appropriate arrangements for reasonable accommodations, which can be found on the web at: <http://www.buffalo.edu/studentlife/who-we-are/departments/accessibility.html>.

Counseling Services. As a student you may experience a range of issues that can cause barriers to learning or reduce your ability to participate in daily activities. These might include strained relationships, anxiety, high levels of stress, alcohol/drug problems, feeling down, health concerns, or unwanted sexual experiences. Counseling, Health Services, and Health Promotion are here to help with these or other concerns. You learn can more about these programs and services by contacting:

Counseling Services: 120 Richmond Quad (North Campus), phone 716-645-2720

202 Michael Hall (South Campus), phone: 716-829-5800

Health Services: Michael Hall (South Campus), phone: 716- 829-3316

Health Promotion: 114 Student Union (North Campus), phone: 716- 645-2837

Sexual Violence. UB is committed to providing a safe learning environment free of all forms of discrimination and sexual harassment, including sexual assault, domestic and dating violence and stalking. If you have experienced gender-based violence (intimate partner violence, attempted or completed sexual assault, harassment, coercion, stalking, etc.), UB has resources to help. This includes academic accommodations, health and counseling services, housing accommodations, helping with legal protective orders, and assistance with reporting the incident to police or other UB officials if you so choose. Please contact UB's Title IX Coordinator at 716-645-2266 for more information. For confidential assistance, you may also contact a Crisis Services Campus Advocate at 716-796-4399.