

CS59200 BDS - New Trends in Big Data Systems

Instructor: Walid G. Aref

Semester: Fall 2022

Class Description: In this course, we will explore recent trends and advances in Big Data Systems. It is assumed that the student has had at least one undergraduate or graduate database systems course. This course is project-centric. Each student will choose a semester-long project. Students can work alone or in groups of two. Students are expected to report regularly during class time and in separate meetings their progress in their project along with presenting research papers that directly contribute to their project. Each student is expected to present 2-3 research presentations each semester. For project reporting, each group will report at least once every two weeks, and for students electing to go faster track, they can report once a week.

We will study topics related to new memory technologies and their impact on database engines, e.g., main-memory database techniques, persistent memory hierarchies and SSDs, their effect on query processing, data indexing, concurrency and lock-free protocols, reliability, and recovery. We will study challenges in realizing NoSQL databases, implications of new hardware technologies on database engine realization including cache hierarchies, NUMA awareness, and vectorization techniques.

Evaluation: Students will be evaluated based on their project progress and accomplishments as well as their class presentations.

- Project (70%): Distributed as follows:
 - (35%) Regular bi-weekly progress using ppt presentations
 - (20%) Final Project Presentation
 - (15%) Final Project Report + Documented Source Code (if applicable)

- Class Presentations (30%)
 - (10% each) 2-3 in-class/zoom presentations
 - (10%) in-class discussions