



Department of Computer Science

## CS59200: Disaggregated Database Systems

Spring 2023

---

<b>Instructor:</b>	Jianguo Wang	<b>Time:</b>	TTh 4:30p-5:45p
<b>E-mail:</b>	<a href="mailto:csjgwang@purdue.edu">csjgwang@purdue.edu</a>	<b>Room:</b>	GRIS 133
<b>Website:</b>	<a href="http://www.cs.purdue.edu/homes/csjgwang/CS590DisaggregatedDB/">www.cs.purdue.edu/homes/csjgwang/CS590DisaggregatedDB/</a>		

---

**Course Description:** Database systems are ubiquitous and fundamental to our information society by managing large-scale data efficiently. Since the inception in the 1970s, database systems are optimized for the monolithic server architecture that physically includes CPU, memory, and storage into the same server box. However, this is completely changed with the emerging trend of resource disaggregation in modern data centers, especially cloud data centers. Resource disaggregation separates CPU, memory, and disk into disaggregated components connected by a fast network fabric, e.g., RDMA. This brings in many benefits including higher resource utilization, better elasticity, and lower cost.

In this course, we will explore the implications of resource disaggregation to building scalable database systems. We will study storage disaggregation, memory disaggregation, and non-volatile memory disaggregation to OLAP/OLTP/HTAP databases. Students are expected to read recent papers from top-tier venues, and work on a semester-long research project on disaggregated database systems. Students can partner in a group (of size up to 3) depending on the scope of the project and the prior approval of the instructor. The lecture is a combination of instructor presentation and student presentation.

This course requires basic knowledge on database systems. Coverage of CS448 or CS541 or CS542 or equivalent courses would be sufficient.

**Pre-requisites:** Database courses such as CS448, CS541, or CS542.

**Textbooks:** No required textbooks. A few optional textbooks:

- Database System Concepts (7th edition). By Avi Silberschatz, Henry F. Korth, and S. Sudarshan.
- Database Management Systems (3rd edition). By Raghu Ramakrishnan and Johannes Gehrke.
- Database Systems: The Complete Book. By Hector Garcia-Molina, Jeffrey Ullman, and Jennifer Widom.

### **Grading Policy:**

- Class participation: 10%
- Paper reading and review: 15%
- Paper presentation: 25%
- Project: 50%
- No midterm and final exam.

**MS/PhD Study:** The course is intended to be included in both MS and PhD study. It contains a semester-long project that takes 50% of the total grade.

### **Tentative Schedule:**

- Week 1: Course Introduction & Projects Overview
- Week 2 & Week 3: OLTP Databases for Storage Disaggregation
- Week 4 & Week 5: OLAP Databases for Storage Disaggregation
- Week 6: HTAP Databases for Storage Disaggregation
- Week 7: Distributed Shared Storage
- Week 8 & 9: Databases for Memory Disaggregation
- Week 10: Distributed Shared Memory
- Week 11: Databases for Non-volatile Memory Disaggregation
- Week 12 & Week 13: CXL to Disaggregated Databases
- Week 14 & Week 15: Final Project Presentation