



*Saylor Academy awards*  
DAN RIMNICEANU

*this certificate for the prescribed program of study for the*  
COMPUTER SCIENCE CURRICULUM

Issue Date: 30 mai 2018

Certificate ID: 11589059



Sean Connor  
Director of Student Affairs  
Saylor Academy



*Saylor Academy awards*

**Dan Rinniceanu**

*this certificate of achievement for*  
**CS410: Advanced Databases**

1 aprilie 2018

Issue Date



11407205

Certificate ID

# CS410: Advanced Databases

## Course Introduction

While [CS403: Introduction to Modern Database Systems](#) covered many of the core concepts behind database management systems, there are many other considerations that should be addressed if you intend to pursue a career in this field. This course will expand upon what you learned about SQL in CS403 and introduce various other advanced topics, including query optimization, concurrency, data warehouses, object-oriented extensions, and XML. While CS403 introduced the basics of database management systems, the additional topics covered in this course will help you become more proficient in writing queries and will expand your knowledge base so that you have a better understanding of the field. By the end of this course, you should have a solid grasp on data warehouses and XML, which will prove to be invaluable as you progress further in your Computer Science studies.

## Unit 1: Advanced SQL

In this unit, we will look at SQL DML commands beyond the basic "select", "join", and "group by", as you will sometimes find it useful to filter a group using a having clause and/or to perform subqueries, which can be used to compare two results set in special ways. We will also study the self-join as well as several set-theoretic operators.

**Completing this unit should take you approximately 13 hours.**

- [Unit 1 Learning OutcomesPage](#)
- **1.1: Relational Algebra Review and Join Commands**

---

  -  [Gonzaga University: Shawn Bowers' "Relational Algebra and Complex SQL"URL](#)  
Read slides 3 through 33 (pages 2-17). As you read this lecture be sure to think of how the necessary mathematical operations are likely to affect the outcome of your database actions.
- **1.2: Additional Join Operations**

---

  -  [Gonzaga University: Shawn Bowers' "Relational Algebra and Complex SQL"URL](#)  
Read slides 3 through 48 (pages 3-24).
- **1.2.1: SELF Join**
  - [TutorialsPoint: "SQL SELF JOINS"URL](#)  
Read this page for a solid overview of the SELF JOINS command and how it is used to join a table. Practice utilizing the SQL commands using MySQL or another available SQL database.

- **1.2.2: FULL Joins**

- [TutorialsPoint: "SQL FULL JOINS"URL](#)

Read this page for a solid overview of the FULL JOINS command and how it is used to achieve a full join of the left and right outer joins of a table. Practice utilizing the SQL commands using MySQL or another available SQL database.

- **1.3: Set-Theoretic Operators**

---

-  [Tom Jewett's "Database Design with UML and SQL: Union and Minus"URL](#)

Read this page for an overview of using the Union, Union All, Minus, and Intersect operators in SQL.

---

- **1.4: The HAVING Clause**

---

-  [Gonzaga University: Shawn Bowers' "Complex SQL"URL](#)

Read slides 3 through 31 (pages 3-16).

---

- [TutorialsPoint: "SQL HAVING Clause"URL](#)

Read this page for a solid overview of the basics of using the HAVING clause command in SQL. Practice utilizing the SQL commands using MySQL or another available SQL database.

---

- **1.5: Views**

-  [Tom Jewett's "Database Design with UML and SQL: SQL Technique: Views and Indexes"URL](#)

Read this page for an overview of views and indexes in SQL.

## **Unit 2: Stored Procedures and Triggers**

Much of the work of database retrieval can be done on the server side as opposed to the client side. The server can execute commands more efficiently and is capable of sending only the result set back to the client, rather than sending the entire data set back and forth between client and server. In this unit, we will look at stored procedures, or blocks of code stored and executed on the server. We will also look at triggers, which are blocks of code that execute on the server when certain events occur, such as the addition or deletion of data from a table. Lastly, we will study rights when running a stored procedure.

**Completing this unit should take you approximately 9 hours.**

- [Unit 2 Learning OutcomesPage](#)

- **2.1: Stored Procedures and Their Usage in Relational Database Management Systems**

---

-  College of San Mateo: Bob Timlin's "Prepared Statements and Store Procedures and Functions"[URL](#)

---

Read this chapter, which will provide and appropriate overview of using Stored Procedures and Functions in MySQL Server.

---

- [MySQL 5.1 Reference Manual: "Stored Procedures and the Grant Table"URL](#)

---

Read sections 18.1 and 18.2 for an overview of stored procedures and stored procedure syntax in RDMS. As you read please keep in mind the effects that stored procedures are likely to have on the Database System.

---

## • 2.2: Stored Functions

---

- [MySQL 5.1 Reference Manual: "Stored Procedures, Functions, Triggers, and LAST\\_INSERT\\_ID\(\)"URL](#)

---

Read sections 18.3 and 18.4 for an overview of stored functions, triggers, and definer and invoker rights in RDMS. As you read compare and contrast the results that can be obtained utilizing Stored Functions in a Database System.

---

## • 2.3: Constraints and Triggers

- [Nguyen Kim Anh's "Structured Query Language"URL](#)

Read this article, which gives an overview of SQL, including constraints in SQL, which are used to ensure database data integrity.

-  [w3resource: "SQL Constraints"URL](#)

Read this article, which describes SQL syntax and examples for constraints.

- [Oracle Community Wiki: "Triggers Tutorial"URL](#)

Read this article, which illustrates constraints and triggers supported by Oracle. A trigger consists of code that is executed when a predefined event occurs at the database, schema, view, or table level. Constraints are simpler than triggers, and are defined at the column or table level, and only check for data integrity.

- [Oracle Community Wiki: "Trigger Development Example"URL](#)

Read this example of a trigger.

## Unit 3: Query Optimization

Queries written in SQL do not always run efficiently. However, there are a number of techniques that you can use to optimize them. In this unit, we will first look at indexes. When tables are indexed on certain fields, the relational database management system can take advantage of quicker retrieval time when searching through the table. Once you have a firm understanding of indexes, please move on to the discussion of query optimization.

**Completing this unit should take you approximately 13 hours.**

- [Unit 3 Learning Outcomes](#)Page

- **3.1: Stages in Query Processing**

---

- [Indian Institute of Technology Bangalore: S. Srinath's "Query Processing and Optimization I"](#)Page
- 

Watch this lecture for a introduction to query processing and optimization. As you watch, be sure to take notes outlining the differences between the various aspects of the Query Plan Language. Determine the strengths and weaknesses between the Logical and Physical Query Plan and languages.

---

- **3.2: Query Processing Algorithms**

---

- [Indian Institute of Technology Bangalore: S. Srinath's "Query Processing and Optimization II"](#)Page
- 

Watch this lecture for a introduction to query processing and optimization with a focus on various algorithms that can be used. As you watch, be sure to understand the differences between the various query processing and optimizing algorithms. Determine which scenarios would be optimal for using a single-pass versus a multi-pass algorithm, etc.

---

- **3.3: Query Plan Execution**

---

- [Indian Institute of Technology Bangalore: S. Srinath's "Query Processing and Optimization III"](#)Page
- 

Watch this lecture for a introduction to query processing and optimization with a focus Query plan execution.

---

- **3.4: Cost-Based Query Optimization**

-  [Gonzaga University: Shawn Bowers' "Query Optimization"](#)URL

Read slides 3 through 20 (pages 2-10) for an overview of cost-based query optimization techniques. As you read this lecture be sure to understand the importance of utilizing the statistical methods for query optimization.

## **Unit 4: Concurrency and Recovery**

Database management systems must be able to support concurrency; that is, they must be able to support multiple users and processes accessing the same records in a table. Modern database management systems handle concurrency in a variety of ways, including through the use of locking and versioning mechanisms. Database management systems must also be able to ensure transaction consistency and recover from catastrophic failures. In this unit, we will look at methods of ensuring that systems are capable of concurrency and recovery.

**Completing this unit should take you approximately 13 hours.**

- [Unit 4 Learning Outcomes](#)Page

- **4.1: Transactions and the ACID Property of Transactions**

---

-  [Avi Silberschatz's "Database System Concepts: Transactions"](#)URL

Read this presentation, which covers transaction properties and serializability.

---

-  [University of the Pacific: Michael Doherty's "Transactions and Serializability"](#)URL

Read this lecture to learn about transaction properties and serializability.

---

- **4.2: Serializability and the Serializability Theorem**

---

- [Carnegie Mellon University: Gregory Kesden's "Serialization"](#)URL

Read these notes on serializability. Be sure that you understand the subcomponents that make up the serializability theorem. Determine the factors that are likely to affect the outcome of the schedule.

---

- **4.3: Two-Phase Locking**

---

-  [Avi Silberschatz's "Database System Concepts: Concurrency Control"](#)URL

Read these slides, which cover two-phase locking timestamps, deadlocks, and multiversion concurrency control. As you read, focus on the mechanism and locks that affect concurrency control.

---

- **4.4: Recovery**

---

-  [Avi Silberschatz's "Database System Concepts: Recovery System"](#)URL

Read these slides, which cover two-phase locking timestamps, and deadlocks. As you read, focus on the components that are likely to affect the recovery of a database systems, as well as types of failures that can occur and how they can be corrected.

## **Unit 5: Database System Architectures**

Databases often need to be distributed across a number of different locations, especially because today's global companies need to be able to access data in multiple locations. While one copy of a database could reside in a location and be accessed remotely, it is often more efficient and reliable to replicate data across servers or to have portions of data on databases that share access across locations. In this unit, we will study distributed databases and related issues.

**Completing this unit should take you approximately 17 hours.**

- [Unit 5 Learning OutcomesPage](#)

- **5.1: Centralized and Client-Server Systems**

---

-  [Avi Silberschatz's "Database System Concepts: Database System Architectures"URL](#)
- 

Read this lecture, which covers server-system architectures and client-server architectures. While you read, compare and contrast the differences between the centralized and client-server systems.

---

- **5.2: Parallel Databases**

---

-  [Avi Silberschatz's "Database System Concepts: Parallel Databases"URL](#)
- 

Read this lecture, which covers parallel databases. While you read, be sure to compare and contrast the components available in parallel databases.

---

- **5.3: Distributed Databases**

-  [Avi Silberschatz's "Database System Concepts: Distributed Databases"URL](#)

Read this lecture, which covers distributed databases.

- **5.3.1: Heterogeneous and Homogeneous Databases**

-  [Massachusetts Institute of Technology: Stuart Madnick's "Distributed Homogeneous Databases"URL](#)

Read this lecture, which covers distributed homogeneous databases.

-  [Massachusetts Institute of Technology: Stuart Madnick's "Distributed Heterogeneous Databases"URL](#)

Read this lecture, which covers distributed heterogeneous databases.

- **5.3.2: Concurrency Control**

- [Indian Institute of Technology Madras: D. Janakiram's "Concurrency Control for Distributed Transaction"Page](#)

Watch this lecture for an introduction to concurrency control mechanisms in distributed databases.

- **5.3.3: Distributed Query Processing**

-  [Free University of Bozen-Bolzano: Johann Gamper's "Distributed Query Processing"URL](#)

Read this lecture, which discusses query processing.

## Unit 6: Data Warehousing

Traditionally, database management systems have been used for transaction processing. Recently, however, the field has seen a trend of creating separate data stores for the purposes of reporting and tracking historical data. In this unit, you will learn about data warehouses and how they are implemented.

**Completing this unit should take you approximately 12 hours.**

- [Unit 6 Learning OutcomesPage](#)

### • 6.1: Data Warehouse Introduction

---

-  [University of Zurich: Michael Böhlen's "Data Warehousing"URL](#)

Read these slides for an introduction to data warehousing and data mining.

---

### • 6.2: SQL OLAP Extensions

---

-  [University of Zurich: Michael Böhlen's "SQL OLAP extensions"URL](#)

Read these slides for an introduction to SQL usage and OLAP extension operations.

---

### • 6.3: An Algebraic OLAP Operator

-  [University of Zurich: Michael Böhlen's "Generalized Multi-Dimensional Join"URL](#)

Read these slides for an introduction to concepts involved in the generalized multi-dimensional join. While you read, focus on the components used for creating the Algebraic OLAP Operator and how it can affect performance.

## Unit 7: Object-Oriented and Object-Relational Databases

As object-oriented programming gained popularity in the 1990s, object-oriented database management systems that incorporated all of the principles of object-oriented design began to emerge. A short time later, a hybrid approach known as object-relational databases was adopted. In this unit, you will first learn about object-oriented databases and then object-relational extensions to relational database management systems.

**Completing this unit should take you approximately 12 hours.**

- [Unit 7 Learning OutcomesPage](#)

### • 7.1: Object-Oriented Data Model

---

- [University of Alberta: Osmar Zaïane's "Object-Oriented Databases"URL](#)

Read these slides for an introduction to concepts related to object-oriented databases. While you read, compare the components available in object-oriented

databases and how they are similar to object-oriented components in programming languages.

---

- [Indian Institute of Technology Bangalore: S. Srinath's "Object Oriented Databases I"Page](#)
- 

Watch this lecture for an introduction to object-oriented databases.

---

- **7.2: Object-Relational Database Systems**

-  [University of California, Santa Cruz: M. Keller's "Object-Relational Databases"URL](#)

Read these slides for an introduction to concepts related to object-relational databases.

- [Indian Institute of Technology Bangalore: S. Srinath's "Object Oriented Databases II"Page](#)

Watch this lecture for an introduction to object-oriented databases.

### **Unit 8: XML and Databases**

While HTML is used to display data on webpages, XML is used to describe data and has quickly become a standard for transmitting data between organizations. In this unit, we will discuss the structure of XML files, how to write queries against such files, and how to integrate XML into a relational database management system.

**Completing this unit should take you approximately 8 hours.**

- [Unit 8 Learning OutcomesPage](#)

- **8.1: XML**

---

-  [Duke University: Jun Yang's "XML Basics"URL](#)
- 

Read these slides for an introduction to XML concepts. While you read, focus on the purpose of using XML and how it differs from other options.

---

- [Indian Institute of Technology Kharagpur: S. Srinath's "Extensible Markup Language \(XML\)"Page](#)
- 

Watch this lecture, which discusses XML and how it is used, including formatting, its use in browsers, its elements, and its attributes.

---

- **8.2: XPath and XQuery**

---

-  [Duke University: Jun Yang's "XPath, XQuery"URL](#)
-

Read these slides for an introduction to concepts related to XPath and XQuery. While you read, focus on the differences between using XPath and other technologies such as SQL.

---

- **8.3: XSLT**

---

-  [Duke University: Jun Yang's "XSLT"URL](#)
- 

Read these slides for an introduction to concepts related to XPath and XQuery.

---

- **8.4: Integrating XML and Databases**

- [Indian Institute of Technology Bangalore: S. Srinath's "XML Databases"Page](#)

Watch this lecture for an introduction to object-oriented databases.