

Answer Key

Name of Examination		Final Assessment Test, Winter 2022-23 Semester, (APRIL 2023)			
Slot: A2 +TA2		Course Mode: CBL		Class Number (s): VL2022230504178	
Course Code:	CSE2031	Course Code:	Principles of Database Management Systems		
Emp. No.:	18844	Faculty Name:	Dr.V.Ilayaraja	School: SCOPE	
Contact No.:	9843189648	Email:	ilayaraja.v@vit.ac.in		

General Instructions (if any): Answer All Questions

Q. No.	Sub-division	Question Text	Marks	Unit / Module No.	Difficulty Level E/A/T	BL	CO
--------	--------------	---------------	-------	-------------------	------------------------	----	----

Answer any TEN Total Marks: 10 X 10 = 100 Marks

1.	a)	<p>List four significant differences between file-processing system and DBMS.</p> <p>1.1 List four significant differences between a file-processing system and a DBMS. Answer: Some main differences between a database management system and a file-processing system are:</p> <ul style="list-style-type: none"> • Both systems contain a collection of data and a set of programs which access that data. A database management system coordinates both the physical and the logical access to the data, whereas a file-processing system coordinates only the physical access. • A database management system reduces the amount of data duplication by ensuring that a physical piece of data is available to all programs authorized to have access to it, whereas data written by one program in a file-processing system may not be readable by another program. • A database management system is designed to allow flexible access to data (i.e., queries), whereas a file-processing system is designed to allow pre-determined access to data (i.e., compiled programs). • A database management system is designed to coordinate multiple users accessing the same data at the same time. A file-processing system is usually designed to allow one or more programs to access different data files at the same time. In a file-processing system, a file can be accessed by two programs concurrently only if both programs have read-only access to the file. 	3	1	E	BL2	CO1
	b)	<p>Describe the two and three tier client-server architecture with neat diagrams in detail.</p> <p style="text-align: center;">Client/Server Architectures</p> <div style="display: flex; justify-content: space-around; align-items: center;"> <div style="text-align: center;"> <p>(a) Two-tier architecture</p> </div> <div style="text-align: center;"> <p>(b) Three-tier architecture</p> </div> </div>	7	1	E	BL1	CO1

2.	a)	<p>Construct an E-R diagram for a car-insurance company whose customers own one or more cars each. Each car has associated with it zero to any number of recorded accidents. Also, indicate the key and mapping constraints.</p> <p>Figure 2.1 E-R diagram for a Car-insurance company.</p>	3	2	A	BL3	CO2
	b)	<p>A university registrar's office maintains data about the following entities:</p> <ul style="list-style-type: none"> (a) courses, including number, title, credits, syllabus, and prerequisites; (b) course offerings, including course number, year, semester, section number, instructor(s), timings, and classroom; (c) students, including student-id, name, and program; and (d) instructors, including identification number, name, department, and title. <p>Further, the enrollment of students in courses and grades awarded to students in each course they are enrolled for must be appropriately modeled.</p> <p>Construct an E-R diagram for the registrar's office. Document all assumptions that you make about the mapping constraints. Also indicate the key for each entity.</p>	7	2	A	BL3	CO2

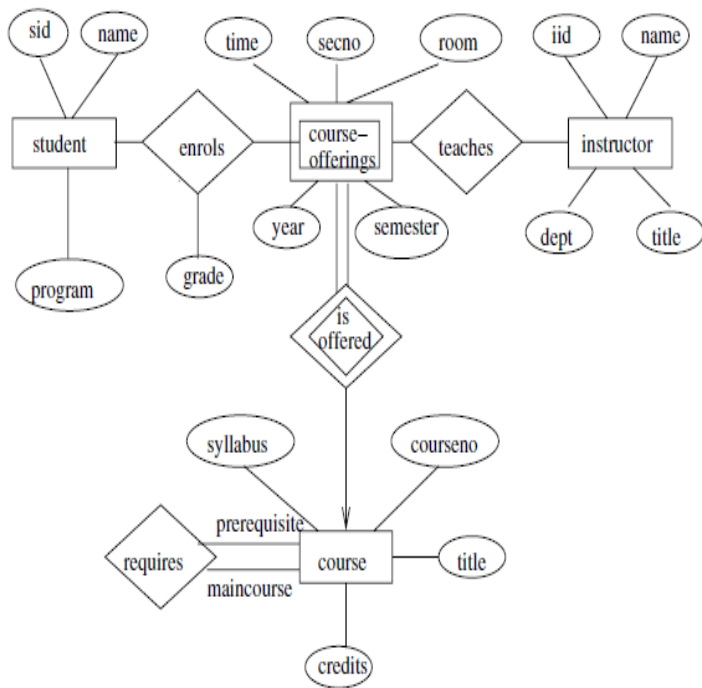



Figure 2.3 E-R diagram for a university.

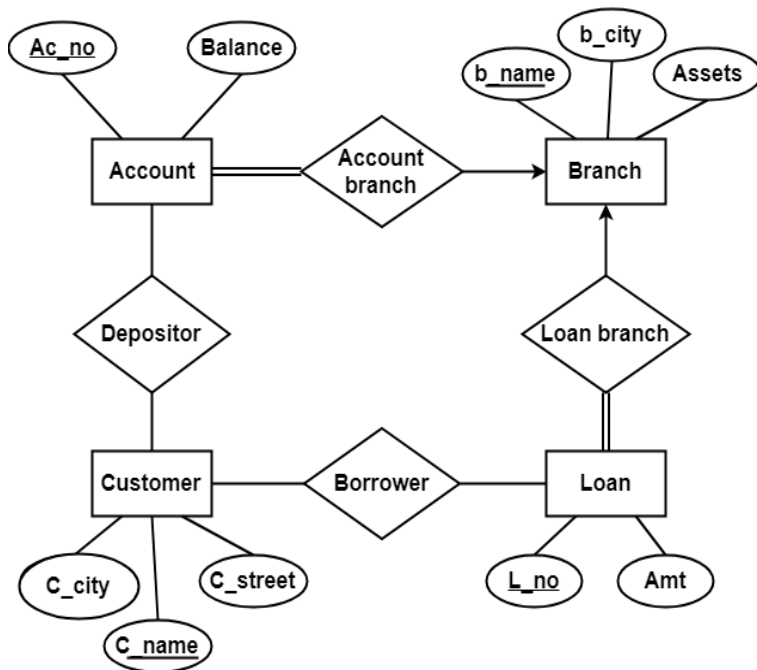
3. a) Design a relational schema corresponding to the given E-R diagram. Also, choose a primary key for a relational schema with respect to mapping cardinality constraints.



Applying the rules, minimum 3 tables will be required-

- MR1 (M1 , M2 , M3 , P1)
- P (P1 , P2)
- NR2 (P1 , N1 , N2)

b) Find the minimum number of tables required to represent the given ER diagram in relational model. Also, choose a primary key for a relational schema with respect to mapping cardinality constraints.



Applying the rules that we have learnt, minimum 6 tables will be required-

- Account (Ac_no , Balance , b_name)
- Branch (b_name , b_city , Assets)
- Loan (L_no , Amt , b_name)
- Borrower (C_name , L_no)
- Customer (C_name , C_street , C_city)
- Depositor (C_name , Ac_no)

4.

The Gill Art Gallery wishes to maintain data on their customers, artists and arts. They may have several arts by each artist in the gallery at one time. The gallery may sell an art, then buy it back at a later date and sell it to another customer.

Consider the following customer history form and Demonstrate the unnormalized relation and convert into 1NF, 2NF and 3NF. Justify each normal form.

Gallery Customer History Form

Customer Name

Jackson, Elizabeth Phone (206) 284-6783
 123 - 4th Avenue
 Fonthill, ON
 L3J 4S4

Purchases Made

Artist	Art Title	Purchase Date	Sales Price
03 - Carol Channing	Laugh with Teeth	09/17/2000	7000.00
15 - Dennis Frings	South toward Emerald Sea	05/11/2000	1800.00
03 - Carol Channing	At the Movies	02/14/2002	5550.00
15 - Dennis Frings	South toward Emerald Sea	07/15/2003	2200.00

UNF:

customer [custno, cust_name, cust_addr, cust_phone, (artist_id, artist_name, art_title, pur_date, price)]

1NF:

customer [custno, cust_name, cust_addr, cust_phone]
 cust_art [custno, art_code, pur_date, artist_id, artist_name, art_title, price]

note: the key chosen for the repeating group is the piece of art itself (a code was assigned), however because a piece of art may be bought by a customer more than once, the purchase date was added as part of the key to make the rows unique.

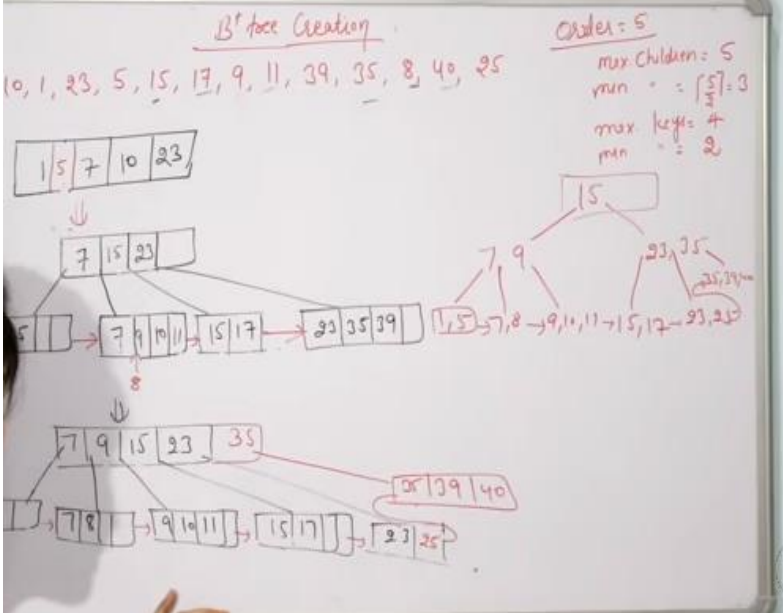
2NF:

customer [custno, cust_name, cust_addr, cust_phone]
 cust_art [custno, art_code, pur_date, price]
 art [art_code, art_title, artist_id, artist_name]

3NF:

customer [custno, cust_name, cust_street, cust_city, cust_prov, cust_pstlcd, cust_phone]

cust_art [custno, art_code, pur_date, price]
 art [art_code, art_title, artist_id(FK)]
 artist [artist_id, artist_fname, artist_lname]


5.	<p>Construct a B+ tree for following key values using order 5 and show the constructed output for each insertion. 7, 10, 1, 23, 5, 15, 17, 9, 11, 39, 35, 8, 40, 25.</p> 	10	4	T	BL3	CO5
6.	<p>a) Consider the following relational database, where the primary keys are underlined. Give an expression in the relational algebra to express each of the following queries:</p> <p>employee (<u>person-name</u>, street, city) works (<u>person-name</u>, company-name, salary) company (<u>company-name</u>, city) manages (<u>person-name</u>, manager-name)</p> <p>i) Find the names of all employees who work for First Bank Corporation. ii) Find the names, street and cities of residence of all employees who work for First Bank Corporation and earn more than Rs.10,000 per annum.</p> <p>a. $\Pi_{\text{person-name}} (\sigma_{\text{company-name} = \text{"First Bank Corporation"}} (\text{works}))$</p> <p>c. $\Pi_{\text{person-name, street, city}} (\sigma_{(\text{company-name} = \text{"First Bank Corporation"} \wedge \text{salary} > 10000)} (\text{works} \bowtie \text{employee}))$</p>	3	5	A	BL3	CO3
	<p>b) Brief about the cost based and heuristic query optimization with necessary diagrams. Cost based: The procedure for generating equivalent expressions can be modified to generate all possible evaluation plans as follows: A new class of equivalence rules, called physical equivalence rules, is added that allows a logical operation, such as a join, to be transformed to a physical operation, such as a hash join, or a nested-loops join. To make the approach work efficiently requires the following: 1. A space-efficient representation of expressions that avoids making multiple copies of the same sub expressions when</p>	7	5	E	BL2	CO3

equivalence rules are applied.


2. Efficient techniques for detecting duplicate derivations of the same expression.
3. A form of dynamic programming based on **memoization**, which stores the optimal query evaluation plan for a sub expression when it is optimized for the first time; subsequent requests to optimize the same sub expression are handled by returning the already memoized plan.
4. Techniques that avoid generating all possible equivalent plans, by keeping track of the cheapest plan generated for any sub expression up to any point of time, and pruning away any plan that is more expensive than the cheapest plan found so far for that sub expression.

Heuristics in Optimization:

- A drawback of cost-based optimization is the cost of optimization itself.
- Although the cost of query optimization can be reduced by clever algorithms, the number of different evaluation plans for a query can be very large, and finding the optimal plan from this set requires a lot of computational effort.
- Hence, optimizers use **heuristics** to reduce the cost of optimization.
- An example of a heuristic rule is the following rule for transforming relational algebra queries:
 - Perform selection operations as early as possible.
 - Perform projections early.



(a) Left-deep join tree



(b) Non-left-deep join tree

**optimization cost budget
plan caching.**

7.	a)	During its execution, a transaction passes through several states, until it finally commits or aborts. List all possible sequences of states through which a transaction may pass. Justify why each state transition may occur.	3	5	A	BL4	CO3
----	----	---	---	---	---	-----	-----

		<p>Answer. The possible sequences of states are:-</p> <p>a. <i>active</i> → <i>partially committed</i> → <i>committed</i>. This is the normal sequence a successful transaction will follow. After executing all its statements it enters the <i>partially committed</i> state. After enough recovery information has been written to disk, the transaction finally enters the <i>committed</i> state.</p> <p>b. <i>active</i> → <i>partially committed</i> → <i>aborted</i>. After executing the last statement of the transaction, it enters the <i>partially committed</i> state. But before enough recovery information is written to disk, a hardware failure may occur destroying the memory contents. In this case the changes which it made to the database are undone, and the transaction enters the <i>aborted</i> state.</p> <p>c. <i>active</i> → <i>failed</i> → <i>aborted</i>. After the transaction starts, if it is discovered at some point that normal execution cannot continue (either due to internal program errors or external errors), it enters the failed state. It is then rolled back, after which it enters the <i>aborted</i> state.</p>														
	b)	<p>What is a cascadeless schedule? Why is cascadeless of schedules desirable? Are there any circumstances under which it would be desirable to allow noncascadeless schedules? Give an example schedule for noncascadeless. Justify your answer.</p> <ul style="list-style-type: none"> • A cascadeless schedule is one where, for each pair of transactions T_i and T_j such that T_j reads data items previously written by T_i, the commit operation of T_i appears before the read operation of T_j. • Cascadeless schedules are desirable because the failure of a transaction does not lead to the aborting of any other transaction. • Of course this comes at the cost of less concurrency. • If failures occur rarely, so that we can pay the price of cascading aborts for the increased concurrency, noncascadeless schedules might be desirable. • noncascadeless shedule <table border="1" style="margin-left: 40px;"> <thead> <tr> <th style="background-color: #ADD8E6;">T_8</th> <th style="background-color: #ADD8E6;">T_9</th> <th style="background-color: #ADD8E6;">T_{10}</th> </tr> </thead> <tbody> <tr> <td>read(A) read(B) write(A)</td> <td>read(A) write(A)</td> <td>read(A)</td> </tr> <tr> <td>abort</td> <td></td> <td></td> </tr> </tbody> </table>	T_8	T_9	T_{10}	read(A) read(B) write(A)	read(A) write(A)	read(A)	abort			7	5	A	BL4	CO3
T_8	T_9	T_{10}														
read(A) read(B) write(A)	read(A) write(A)	read(A)														
abort																
8.	a)	<p>What type of problem is there in a following schedule? Justify the drawback of this problem.</p> <pre> T₃: lock-X(B); read(B); B := B - 50; write(B); lock-X(A); read(A); A := A + 50; write(A); unlock(B); unlock(A). </pre> <p>unlocking is delayed to the end of the transaction.</p>	3	6	A	BL4	CO4									
	b)	<p>Write the rules of timestamp based protocol when transaction issues a read and write operation. Explain in detail.</p>	7	6	A	BL2	CO4									

	<ul style="list-style-type: none"> Suppose a transaction T_i issues a read(Q) <ol style="list-style-type: none"> If $TS(T_i) < W\text{-timestamp}(Q)$, then T_i needs to read a value of Q that was already overwritten. Hence, the read operation is rejected, and T_i is rolled back. If $TS(T_i) \geq W\text{-timestamp}(Q)$, then the read operation is executed, and $R\text{-timestamp}(Q)$ is set to the maximum of $R\text{-timestamp}(Q)$ and $TS(T_i)$. Suppose that transaction T_i issues write(Q). <ol style="list-style-type: none"> If $TS(T_i) < R\text{-timestamp}(Q)$, then the value of Q that T_i is producing was needed previously, and the system assumed that that value would never be produced. Hence, the write operation is rejected, and T_i is rolled back. If $TS(T_i) < W\text{-timestamp}(Q)$, then T_i is attempting to write an obsolete value of Q. Hence, this write operation is rejected, and T_i is rolled back. Otherwise, the write operation is executed, and $W\text{-timestamp}(Q)$ is set to $TS(T_i)$. 																													
9.	<p>Consider the following log instances a, b and c in Figure A and B. If log on stable storage at time of crash at every instance, State the recovery mechanism for each log instance in both deferred (consider Figure A) and immediate (consider Figure B) database modification.</p> <p>Figure A - Deferred database modification.</p> <table border="1" data-bbox="268 1032 1038 1368"> <tr> <td>$\langle T_0 \text{ start} \rangle$</td> <td>$\langle T_0 \text{ start} \rangle$</td> <td>$\langle T_0 \text{ start} \rangle$</td> </tr> <tr> <td>$\langle T_0, A, 950 \rangle$</td> <td>$\langle T_0, A, 950 \rangle$</td> <td>$\langle T_0, A, 950 \rangle$</td> </tr> <tr> <td>$\langle T_0, B, 2050 \rangle$</td> <td>$\langle T_0, B, 2050 \rangle$</td> <td>$\langle T_0, B, 2050 \rangle$</td> </tr> <tr> <td></td> <td>$\langle T_0 \text{ commit} \rangle$</td> <td>$\langle T_0 \text{ commit} \rangle$</td> </tr> <tr> <td></td> <td>$\langle T_1 \text{ start} \rangle$</td> <td>$\langle T_1 \text{ start} \rangle$</td> </tr> <tr> <td></td> <td>$\langle T_1, C, 600 \rangle$</td> <td>$\langle T_1, C, 600 \rangle$</td> </tr> <tr> <td></td> <td></td> <td>$\langle T_1 \text{ commit} \rangle$</td> </tr> <tr> <td>(a)</td> <td>(b)</td> <td>(c)</td> </tr> </table> <p>(a) No redo actions need to be taken</p> <p>(b) redo(T_0) must be performed since $\langle T_0 \text{ commit} \rangle$ is present</p> <p>(c) redo(T_0) must be performed followed by redo(T_1) since $\langle T_0 \text{ commit} \rangle$ and $\langle T_1 \text{ commit} \rangle$ are present</p>	$\langle T_0 \text{ start} \rangle$	$\langle T_0 \text{ start} \rangle$	$\langle T_0 \text{ start} \rangle$	$\langle T_0, A, 950 \rangle$	$\langle T_0, A, 950 \rangle$	$\langle T_0, A, 950 \rangle$	$\langle T_0, B, 2050 \rangle$	$\langle T_0, B, 2050 \rangle$	$\langle T_0, B, 2050 \rangle$		$\langle T_0 \text{ commit} \rangle$	$\langle T_0 \text{ commit} \rangle$		$\langle T_1 \text{ start} \rangle$	$\langle T_1 \text{ start} \rangle$		$\langle T_1, C, 600 \rangle$	$\langle T_1, C, 600 \rangle$			$\langle T_1 \text{ commit} \rangle$	(a)	(b)	(c)	10	6	A	BL4	CO4
$\langle T_0 \text{ start} \rangle$	$\langle T_0 \text{ start} \rangle$	$\langle T_0 \text{ start} \rangle$																												
$\langle T_0, A, 950 \rangle$	$\langle T_0, A, 950 \rangle$	$\langle T_0, A, 950 \rangle$																												
$\langle T_0, B, 2050 \rangle$	$\langle T_0, B, 2050 \rangle$	$\langle T_0, B, 2050 \rangle$																												
	$\langle T_0 \text{ commit} \rangle$	$\langle T_0 \text{ commit} \rangle$																												
	$\langle T_1 \text{ start} \rangle$	$\langle T_1 \text{ start} \rangle$																												
	$\langle T_1, C, 600 \rangle$	$\langle T_1, C, 600 \rangle$																												
		$\langle T_1 \text{ commit} \rangle$																												
(a)	(b)	(c)																												

	<p>Figure B - Immediate database modification</p> <table border="1" style="margin-left: auto; margin-right: auto;"> <tr> <td style="text-align: center;"> $\langle T_0 \text{ start} \rangle$ $\langle T_0, A, 1000, 950 \rangle$ $\langle T_0, B, 2000, 2050 \rangle$ (a) </td> <td style="text-align: center;"> $\langle T_0 \text{ start} \rangle$ $\langle T_0, A, 1000, 950 \rangle$ $\langle T_0, B, 2000, 2050 \rangle$ $\langle T_0 \text{ commit} \rangle$ $\langle T_1 \text{ start} \rangle$ $\langle T_1, C, 700, 600 \rangle$ (b) </td> <td style="text-align: center;"> $\langle T_0 \text{ start} \rangle$ $\langle T_0, A, 1000, 950 \rangle$ $\langle T_0, B, 2000, 2050 \rangle$ $\langle T_0 \text{ commit} \rangle$ $\langle T_1 \text{ start} \rangle$ $\langle T_1, C, 700, 600 \rangle$ $\langle T_1 \text{ commit} \rangle$ (c) </td> </tr> </table> <p>(a) undo (T_0): B is restored to 2000 and A to 1000. (b) undo (T_1) and redo (T_0): C is restored to 700, and then A and B are set to 950 and 2050 respectively. (c) redo (T_0) and redo (T_1): A and B are set to 950 and 2050 respectively. Then C is set to 600</p>	$\langle T_0 \text{ start} \rangle$ $\langle T_0, A, 1000, 950 \rangle$ $\langle T_0, B, 2000, 2050 \rangle$ (a)	$\langle T_0 \text{ start} \rangle$ $\langle T_0, A, 1000, 950 \rangle$ $\langle T_0, B, 2000, 2050 \rangle$ $\langle T_0 \text{ commit} \rangle$ $\langle T_1 \text{ start} \rangle$ $\langle T_1, C, 700, 600 \rangle$ (b)	$\langle T_0 \text{ start} \rangle$ $\langle T_0, A, 1000, 950 \rangle$ $\langle T_0, B, 2000, 2050 \rangle$ $\langle T_0 \text{ commit} \rangle$ $\langle T_1 \text{ start} \rangle$ $\langle T_1, C, 700, 600 \rangle$ $\langle T_1 \text{ commit} \rangle$ (c)					
$\langle T_0 \text{ start} \rangle$ $\langle T_0, A, 1000, 950 \rangle$ $\langle T_0, B, 2000, 2050 \rangle$ (a)	$\langle T_0 \text{ start} \rangle$ $\langle T_0, A, 1000, 950 \rangle$ $\langle T_0, B, 2000, 2050 \rangle$ $\langle T_0 \text{ commit} \rangle$ $\langle T_1 \text{ start} \rangle$ $\langle T_1, C, 700, 600 \rangle$ (b)	$\langle T_0 \text{ start} \rangle$ $\langle T_0, A, 1000, 950 \rangle$ $\langle T_0, B, 2000, 2050 \rangle$ $\langle T_0 \text{ commit} \rangle$ $\langle T_1 \text{ start} \rangle$ $\langle T_1, C, 700, 600 \rangle$ $\langle T_1 \text{ commit} \rangle$ (c)							
10	<p>Describe about the features, advantages and disadvantages of various NoSQL data models. Also give an example for each and brief about an example.</p> <p>Key Value Stores</p> <ul style="list-style-type: none"> • Most Based on Dynamo: Amazon Highly Available Key-Value Store • Data Model: <ul style="list-style-type: none"> – Global key-value mapping – Big scalable HashMap – Highly fault tolerant (typically) • Examples: <ul style="list-style-type: none"> – Redis, Riak, Voldemort Pros: <ul style="list-style-type: none"> – Simple data model – Scalable • Cons <ul style="list-style-type: none"> – Create your own “foreign keys” – Poor for complex data <p>Column Family</p> <ul style="list-style-type: none"> • Most Based on BigTable: Google’s Distributed Storage System for Structured Data • Data Model: <ul style="list-style-type: none"> – A big table, with column families – Map Reduce for querying/processing • Examples: <ul style="list-style-type: none"> – HBase, HyperTable, Cassandra Pros: <ul style="list-style-type: none"> – Supports Simi-Structured Data – Naturally Indexed (columns) – Scalable • Cons <ul style="list-style-type: none"> – Poor for interconnected data <p>Document Databases</p> <p>Data Model:</p> <ul style="list-style-type: none"> – A collection of documents – A document is a key value collection – Index-centric, lots of map-reduce 	10	7	E	BL2	CO6			

	<ul style="list-style-type: none"> • Examples: <ul style="list-style-type: none"> – CouchDB, MongoDB Pros: <ul style="list-style-type: none"> – Simple, powerful data model – Scalable • Cons <ul style="list-style-type: none"> – Poor for interconnected data – Query model limited to keys and indexes – Map reduce for larger queries <p>Graph Databases</p> <p>Data Model:</p> <ul style="list-style-type: none"> – Nodes and Relationships <ul style="list-style-type: none"> • Examples: <ul style="list-style-type: none"> – Neo4j, OrientDB, Infinite Graph, Allegro Graph Pros: <ul style="list-style-type: none"> – Powerful data model, as general as RDBMS – Connected data locally indexed – Easy to query • Cons <ul style="list-style-type: none"> – Sharding (lots of people working on this) • Scales UP reasonably well – Requires rewiring your brain 					
--	--	--	--	--	--	--