

USN 

## Internal Assessment Test 1 – March 2024

Sub:	NoSQL Database				Sub Code:	21CS745	Branch:	AIML	
Date:		Duration:	90 min's	Max Marks:	50	Sem/Sec:	VIII / A		OBE
<u>Answer any FIVE FULL Questions</u>							MARKS	CO	RBT
1 a)	Define materialized view. How are they different from views? Briefly explain the two main strategies to build materialized view					10M	CO1	L2	
2 a)	What is NoSQL? Explain about aggregate data models with neat diagram. Considering example of Relational data models.					10M	CO1	L2	
3 a)	Explain the data management and access in column family data stores with example					6M	CO1	L2	
3 b)	Briefly describe the value of relational databases.					4M	CO1	L2	
4 a)	Identify the type of conflict in the following scenario, How can it be solved by applying suitable technique? Alice and Bob share a common Google sheet online. Both read a file. Alice updates the document and forgets to save the file. On the other hand Bob updates the sheet and saves the file. The content updated by Alice overwritten by Bob. The data updated by Alice is lost.					10M	CO2	L3	
5 a)	Explain about Update Consistency and Read Consistency, with an example.					10M	CO2	L2	

USN 

## Internal Assessment Test 1 – March 2024

Sub:	NoSQL Database				Sub Code:	21CS745	Branch:	AIML	
Date:		Duration:	90 min's	Max Marks:	50	Sem/Sec:	VIII / A		OBE
<u>Answer any FIVE FULL Questions</u>							MARKS	CO	RBT
1 a)	Define materialized view. How are they different from views? Briefly explain the two main strategies to build materialized view					10M	CO1	L2	
2 a)	What is NoSQL? Explain about aggregate data models with neat diagram. Considering example of Relational data models.					10M	CO1	L2	
3 a)	Explain the data management and access in column family data stores with example					6M	CO1	L2	
3 b)	Briefly describe the value of relational databases.					4M	CO1	L2	
4 a)	Identify the type of conflict in the following scenario, How can it be solved by applying suitable technique? Alice and Bob share a common Google sheet online. Both read a file. Alice updates the document and forgets to save the file. On the other hand Bob updates the sheet and saves the file. The content updated by Alice overwritten by Bob. The data updated by Alice is lost.					10M	CO2	L3	
5 a)	Explain about Update Consistency and Read Consistency, with an example.					10M	CO2	L2	

6 a)	Explain briefly impedance mismatch, with a neat diagram.	5M	CO1	L2
6 b)	Use the above diagram and answer the following questions. <div data-bbox="336 235 1059 611" style="border: 1px solid black; padding: 10px; margin: 10px auto; width: fit-content;"> <pre> graph TD     A((A)) -- "sister-in-law of" --&gt; S((S))     A -- "married to" --&gt; C((C))     A -- "listens to" --&gt; RM((Rock Music))     C -- "brother of" --&gt; S     C -- "drives" --&gt; BMW((BMW))     C -- "colleague of" --&gt; E((E))     E -- "works for" --&gt; D((D))     E -- "has a pet" --&gt; Bruno((Bruno))         </pre> </div> <ol style="list-style-type: none"> <li>Who listens to rock music and works for D?</li> <li>Who works for D and has married to colleagues?</li> <li>Who listen to rock music?</li> <li>How are A and S related to each other and also to C?</li> <li>What are the genders of C and A?</li> </ol>	5M	CO1	L3

Faculty Signature

CCI Signature

HOD Signature

6 a)	Explain briefly impedance mismatch, with a neat diagram.	5M	CO1	L2
6 b)	Use the above diagram and answer the following questions. <div data-bbox="400 1267 1123 1644" style="border: 1px solid black; padding: 10px; margin: 10px auto; width: fit-content;"> <pre> graph TD     A((A)) -- "sister-in-law of" --&gt; S((S))     A -- "married to" --&gt; C((C))     A -- "listens to" --&gt; RM((Rock Music))     C -- "brother of" --&gt; S     C -- "drives" --&gt; BMW((BMW))     C -- "colleague of" --&gt; E((E))     E -- "works for" --&gt; D((D))     E -- "has a pet" --&gt; Bruno((Bruno))         </pre> </div> <ol style="list-style-type: none"> <li>Who listens to rock music and works for D?</li> <li>Who works for D and has married to colleagues?</li> <li>Who listen to rock music?</li> <li>How are A and S related to each other and also to C?</li> <li>What are the genders of C and A?</li> </ol>	5M	CO1	L3

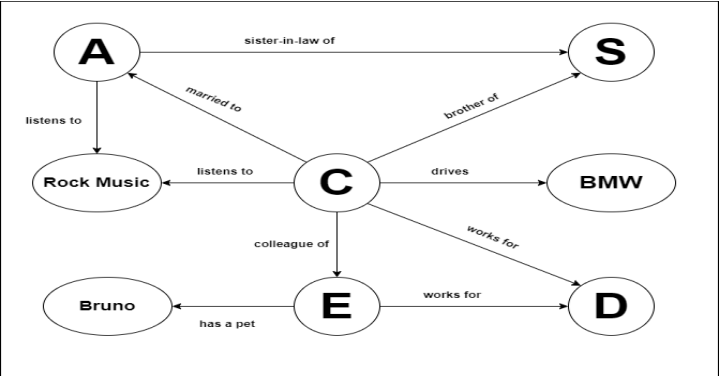
Faculty Signature

CCI Signature

HOD Signature

Scheme of Evaluation  
Internal Assessment Test 1 – March 2024

Sub:	NoSQL Database				Sub Code:	21CS745	Branch:	AIML	
Date:		Duration:	90 min's	Max Marks:	50	Sem/Sec:	VIII / A		OBE
<u>Answer any FIVE FULL Questions</u>									
Q No	Description				Marks Distribution			Max MARKS	
1 a)	Define materialized view. How are they different from views? Briefly explain the two main strategies to build materialized view Define materialized view Two approaches Eager approach Batch approach				2M	4M	4M	10M	10M
2 a)	What is NoSQL? Explain about aggregate data models with neat diagram. Considering example of Relational data models. Definition NoSQL Diagram Explanation				2M	3M	5M	10M	10M
3 a)	Explain the data management and access in column family data stores with example Diagram Explanation				3M	3M		6M	6M
3 b)	Briefly describe the value of relational databases. Getting persistent data Concurrency Integration				2M	1M	1M	4M	4M
4 a)	Identify the type of conflict in the following scenario, How can it be solved? Alice and Bob share a common Google sheet online. Both read a file. Alice updates the document and forgets to save the file. On the other hand Bob updates the sheet and saves the							10M	10M

	file. The content updated by Alice overwritten by Bob. The data updated by Alice is lost. Write-write conflict Locks used explanation with example	2M 4M 4M		
5 a)	Explain about Update Consistency and Read Consistency, with an example. Update Consistency Examples Read Consistency Examples	2M 3M 2M 3M	10M	10M
6 a)	Explain briefly impedance mismatch, with a neat diagram. Diagram explanation	2M 3M	5M	5M
6 b)	Use the above diagram and answer the following questions.  <p>f) Who listens to rock music and works for D? g) Who works for D and has married to colleagues? h) Who listen to rock music? i) How are A and S related to each other and also to C? j) What are the genders of C and A?</p>	1M each	5M	5M

## Solution

**Q. 1 a) Define materialized view. How are they different from views? Briefly explain the two main strategies to build materialized view**

**Ans.**

### **Materialized Views**

- To cope with this, materialized views were invented, which are views that are computed in advance and cached on disk. Materialized views are effective for data that is read heavily but can stand being somewhat stale.
- Although NoSQL databases don't have views, they may have precomputed and cached queries, and they reuse the term "materialized view" to describe them. Often, NoSQL databases create materialized views using a map-reduce computation.

### **There are two strategies to building a materialized view**

- The first is the eager approach, where you update the materialized view at the same time you update the base data for it. In this case, adding an order would also update the purchase history aggregates for each product.
- This approach is good when you have more frequent reads of the materialized view than you have writes, and you want the materialized views to be as fresh as possible. The application database approach is valuable here, as it makes it easier to ensure that any updates to base data also update materialized views.
- If you don't want to pay that overhead on each update, you can run batch jobs to update the materialized views at regular intervals as per requirements.
- This approach is good when you have more frequent reads of the materialized view than you have writes, and you want the materialized views to be as fresh as possible. The application database approach is valuable here, as it makes it easier to ensure that any updates to base data also update materialized views.
- If you don't want to pay that overhead on each update, you can run batch jobs to update the materialized views at regular intervals as per requirements.
- You can build materialized views outside the database by reading the data, computing the view, and saving it back to the database.
- More often, databases will support building materialized views themselves.
- In this case, you provide the computation that needs to be done, and the database executes the computation when needed according to some parameters that you configure. This is particularly handy for eager updates of views with incremental map-reduce.

### **2 a) What is NoSQL? Explain about aggregate data models with neat diagram. Considering example of Relational data models.**

**Ans.** NoSQL, also referred to as "not only SQL" or "non-SQL", is an approach to database design that enables the storage and querying of data outside the traditional structures found in relational databases. Example of Relations and Aggregates At this point, an example may help explain what we're talking about. Let's assume we have to build an e-commerce website; we are going to be selling items directly to customers over the web, and we will have to store information about users, our product catalog, orders, shipping addresses, billing addresses, and payment data. We can use this scenario to model the data using a relation data store as well as NoSQL data stores and talk about their pros and cons. For a relational database, we might start with a data model shown in Figure 2.1. Figure 2.1. Data model oriented around a relational database (using UML notation

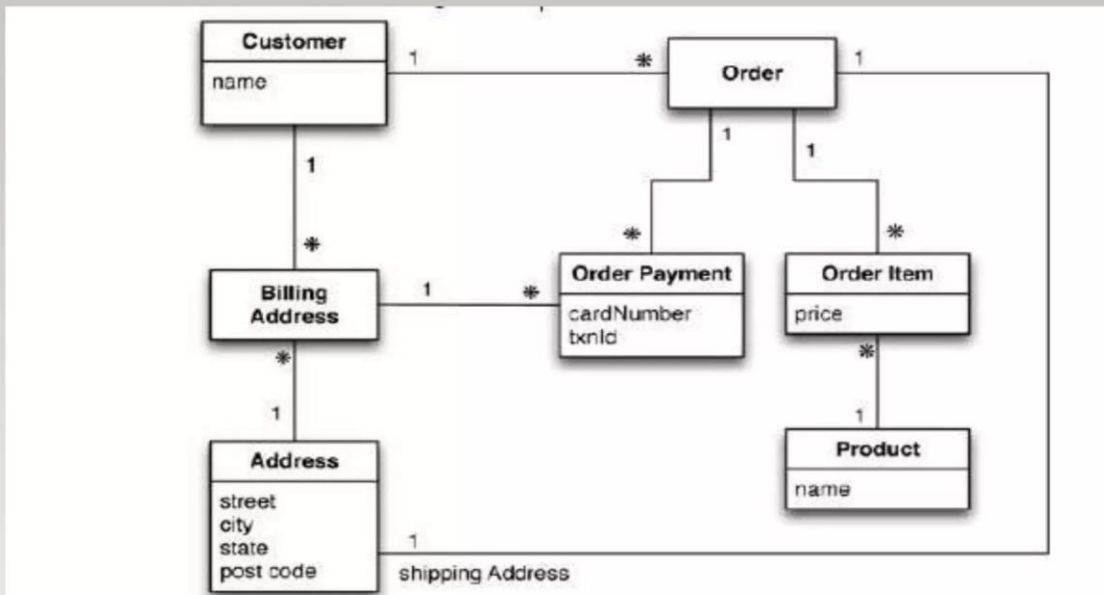


Figure 2.1. Data model oriented around a relational database (using UML notation)

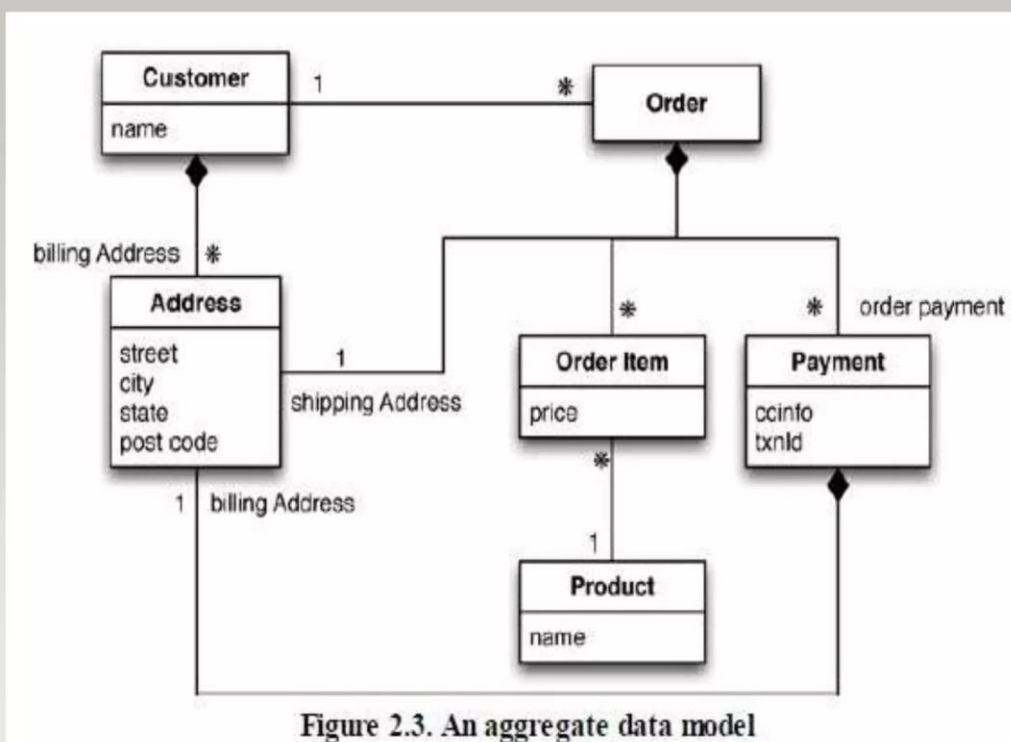


Figure 2.3. An aggregate data model

In this model, we have two main aggregates: customer and order. We've used the black-diamond composition marker in UML to show how data fits into the aggregation structure. The customer contains a list of billing addresses; the order contains a list of order items, a shipping address, and payments. The payment itself contains a billing address for that payment. A single logical address record appears three times in the example data, but instead of using IDs it's treated as a value and copied each time. This fits the domain where we would not want the shipping address, nor the payment's billing address, to change. In a



relational database, we would ensure that the address rows aren't updated for this case, making a new row instead. With aggregates, we can copy the whole address structure into the aggregate as we need to. The link between the customer and the order isn't within either aggregate—it's a relationship between aggregates. Similarly, the link from an order item would cross into a separate aggregate structure for products, which we haven't gone into. We've shown the product name as part of the order item here—this kind of denormalization is similar to the tradeoffs with relational databases, but is more common with aggregates because we want to minimize the number of aggregates we access during a data interaction. The important thing to notice here isn't the particular way we've drawn the aggregate boundary so much as the fact that you have to think about accessing that data—and make that part of your thinking when developing the application data model. Indeed we could draw our aggregate boundaries differently, putting all the orders for a customer into the customer aggregate.

**Q 3 a) Explain the data management and access in column family data stores with example**

**Ans.**

### **Column-Family Stores**

- One of the early and powerful NoSQL databases was Google's BigTable, it is a two-level map. It has been a model that influenced later databases such as HBase and Cassandra.
- These databases with a BigTable-style data model are often referred to as column stores. The thing that made them different was the way in which they physically stored data.
- Most databases have a row as a unit of storage which, in particular, helps write performance. However, there are many scenarios where writes are rare, but you often need to read a few columns of many rows at once.
- In this situation, it's better to store groups of columns for all rows as the basic storage unit—which is why these databases are called column stores.
- BigTable and its next generation follow this notion of storing groups of columns (column families) together, we refer to this as column-family databases.
- Column-family model is a two-level aggregate structure. As with key-value stores, the first key is often described as a row identifier, picking up the aggregate of interest. The difference with column-family structures is that this row aggregate is itself formed of a map of more detailed values. These second-level values are referred to as columns. As well as accessing the row as a whole, operations also allow picking out a particular column, so to get a particular customer's name from you could do something like `get('1234', 'name')`.

Column-family databases organize their columns into column families. Each column has to be part of a single column family, and the column acts as a unit for access, with the assumption that data for a particular column family will usually be accessed together.

- This also gives you a couple of ways to think about how the data is structured.
- Row-oriented: Each row is an aggregate (for example, customer with the ID of 1234) with column families representing useful chunks of data (profile, order history) within that aggregate.
- Column-oriented: Each column family defines a record type (e.g., customer profiles) with rows for each of the records. You then think of a row as the join of records in all column families.

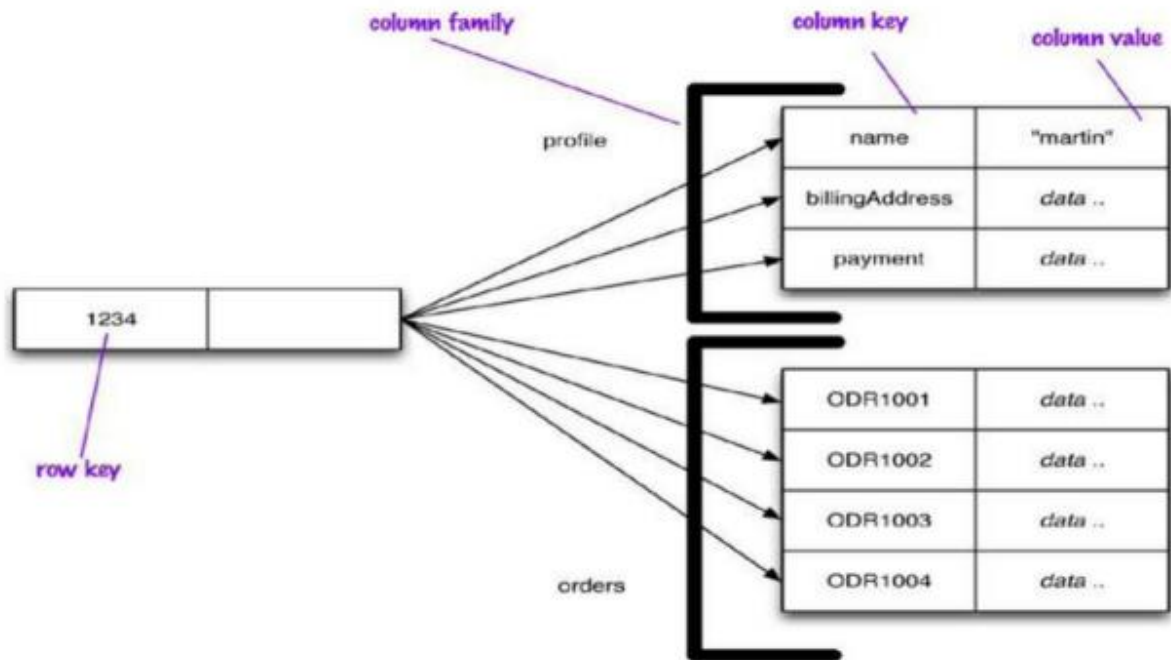


Figure 5: Representing customer info in a column-family structure

- This latter aspect reflects the columnar nature of column-family databases. Since the database knows about these common groupings of data, it can use this information for its storage and access behaviour.
- Cassandra uses the terms “wide” and “skinny.”
- Skinny rows have few columns, with the same columns used across the many different rows.
- In this case, the column family defines a record type, each row is a record, and each column is a field.
- A wide row has many columns (perhaps thousands), with rows having very different columns.
- A wide column family models a list, with each column being one element in that list.



**Q. 3 b) Briefly describe the value of relational databases.**

**Ans.**

## **1.1 The Value of Relational Databases**

Relational databases have become such an embedded part of our computing culture that it's easy to take them for granted. It's therefore useful to revisit the benefits they provide.

### **1.1.1. Getting at Persistent Data**

Probably the most obvious value of a database is keeping large amounts of persistent data. Most computer architectures have the notion of two areas of memory: a fast volatile "main memory" and a larger but slower "backing store." Main memory is both limited in space and loses all data when you lose power or something bad happens to the operating system. Therefore, to keep data around, we write it to a backing store, commonly seen as a disk (although these days that disk can be persistent memory).

The backing store can be organized in all sorts of ways. For many productivity applications (such as word processors), it's a file in the file system of the operating system. For most enterprise applications, however, the backing store is a database. The database allows more flexibility than a file system in storing large amounts of data in a way that allows an application program to get at small bits of that information quickly and easily.

---

### **1.1.2. Concurrency**

Enterprise applications tend to have many people looking at the same body of data at once, possibly modifying that data. Most of the time they are working on different areas of that data, but occasionally they operate on the same bit of data. As a result, we have to worry about coordinating these interactions to avoid such things as double booking of hotel rooms.

Concurrency is notoriously difficult to get right, with all sorts of errors that can trap even the most careful programmers. Since enterprise applications can have lots of users and other systems all working concurrently, there's a lot of room for bad things to happen. Relational databases help handle this by controlling all access to their data through transactions. While this isn't a cure-all (you still have to handle a transactional error when you try to book a room that's just gone), the transactional mechanism has worked well to contain the complexity of concurrency.

Transactions also play a role in error handling. With transactions, you can make a change, and if an error occurs during the processing of the change you can roll back the transaction to clean things up.

### **1.1.3. Integration**

Enterprise applications live in a rich ecosystem that requires multiple applications, written by different teams, to collaborate in order to get things done. This kind of inter-application collaboration is awkward because it means pushing the human organizational boundaries. Applications often need to use the same data and updates made through one application have to be visible to others.

A common way to do this is shared database integration [Hohpe and Woolf] where multiple applications store their data in a single database. Using a single database allows all the applications to use each others' data easily, while the database's concurrency control handles multiple applications in the same way as it handles multiple users in a single application.

### 1.1.4. A (Mostly) Standard Model

Relational databases have succeeded because they provide the core benefits we outlined earlier in a (mostly) standard way. As a result, developers and database professionals can learn the basic relational model and apply it in many projects. Although there are differences between different relational databases, the core mechanisms remain the same: Different vendors' SQL dialects are similar, transactions operate in mostly the same way.

**4 a) Identify the type of conflict in the following scenario, How can it be solved by applying suitable technique? Alice and Bob share a common Google sheet online. Both read a file. Alice updates the document and forgets to save the file. On the other hand Bob updates the sheet and saves the file. The content updated by Alice overwritten by Bob. The data updated by Alice is lost**

**Ans.** This issue is called as write-write conflict: two people updating the same data item at the same time. When the writes reach the server, the server will serialize them—decide to apply one, then the other. In this case Alice's is a lost update. Here the lost update is not a big problem, but often it is. We see this as a failure of consistency because Bob's update was based on the state before Alice's update, yet was applied after it. She forgot to save the file. Approaches for maintaining consistency in the face of concurrency are often described as pessimistic or optimistic. A pessimistic approach works by preventing conflicts from occurring; an optimistic approach lets conflicts occur, but detects them and takes action to sort them out. For update conflicts, the most common pessimistic approach is to have write locks, so that in order to change a value you need to acquire a lock, and the system ensures that only one client can get a lock at a time.

So Alice and Bob would both attempt to acquire the write lock, but only Alice (the first one) would succeed. Bob would then see the result of Bob's write before deciding whether to make his own update. There is another optimistic way to handle a write-write conflict—save both updates and record that they are in conflict. This approach is familiar to many programmers from version control systems, particularly distributed version control systems that by their nature will often have conflicting commits. The next step again follows from version control: You have to merge the two updates somehow. Maybe you show both values to the user and ask them to sort it out—this is what happens if you update the same contact on your phone and your computer. Alternatively, the computer may be able to perform the merge itself; if it was a phone formatting issue, it may be able to realize that and apply the new number with the standard format. Any automated merge of write-write conflicts is highly domain-specific and needs to be programmed for each particular case. Often, when people first encounter these issues, their reaction is to prefer pessimistic concurrency because they are determined to avoid conflicts. While in some cases this is the right answer, there is always a tradeoff.

Concurrent programming involves a fundamental tradeoff between safety (avoiding errors such as update conflicts) and liveness (responding quickly to clients). Pessimistic approaches often severely degrade the responsiveness of a system to the degree that it becomes unfit for its purpose. This problem is made worse by the danger of errors—pessimistic concurrency often leads to deadlocks, which are hard to prevent and debug. Replication makes it much more likely to run into write-write conflicts. If different nodes have different copies of some data which can be independently updated, then you'll get conflicts unless you take specific measures to avoid them. Using a single node as the target for all writes for some data makes it much easier to maintain update consistency. Of the distribution models we discussed earlier, all but peer-to-peer replication do this.

**5 a) Explain about Update Consistency and Read Consistency, with an example.**

**Ans.**

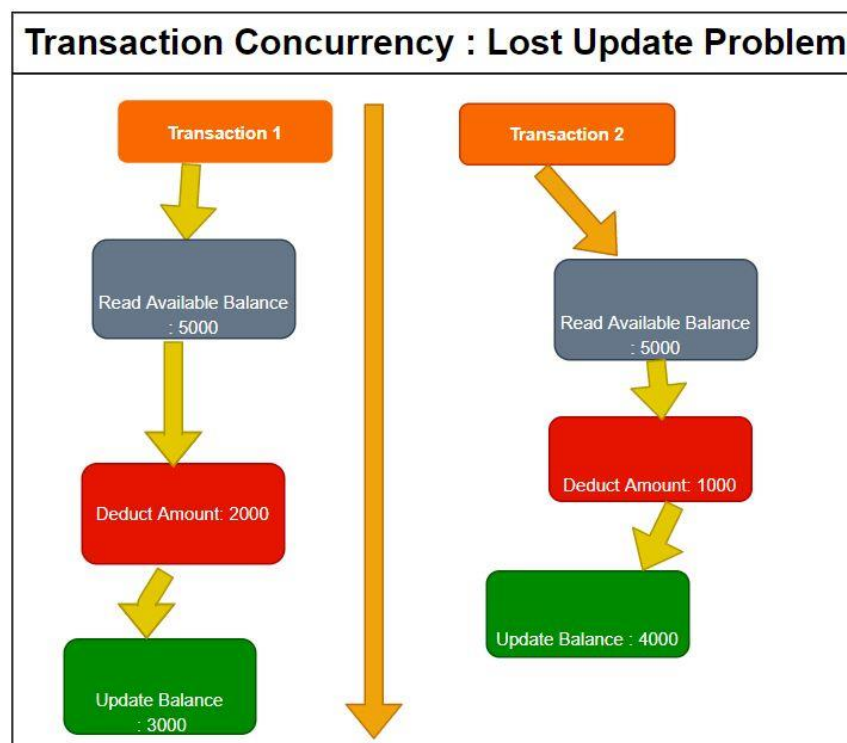


## 2.1 Update Consistency

We'll begin by considering updating a telephone number. Coincidentally, Martin and Pramod are looking at the company website and notice that the phone number is out of date. Implausibly, they both have update access, so they both go in at the same time to update the number. To make the example interesting, we'll assume they update it slightly differently, because each uses a slightly different format. This issue is called a **write-write conflict**: two people updating the same data item at the same time.

When the writes reach the server, the server will **serialize** them—decide to apply one, then the other. Let's assume it uses alphabetical order and picks Martin's update first, then Pramod's. Without any concurrency control, Martin's update would be applied and immediately overwritten by Pramod's. In this case Martin's is a **lost update**. Here the lost update is not a big problem, but often it is. We see this as a failure of consistency because Pramod's update was based on the state before Martin's update, yet was applied after it.

Approaches for maintaining consistency in the face of concurrency are often described as pessimistic or optimistic. A **pessimistic** approach works by preventing conflicts from occurring; an **optimistic** approach lets conflicts occur, but detects them and takes action to sort them out. For update conflicts, the most common pessimistic approach is to have write locks, so that in order to



## Read Consistency

Having a data store that maintains update consistency is one thing, but it doesn't guarantee that readers of that data store will always get consistent responses to their requests. Let's imagine we have an order with line items and a shipping charge. The shipping charge is calculated based on the line items in the order. If we add a line item, we thus also need to recalculate and update the shipping charge. In a relational database, the shipping charge and line items will be in separate tables. The danger of inconsistency is that Martin adds

a line item to his order, Pramod then reads the line items and shipping charge, and then Martin updates the shipping charge. This is an inconsistent read or read-write conflict: In Figure 2.1 Pramod has done a read in the middle of Martin's write.

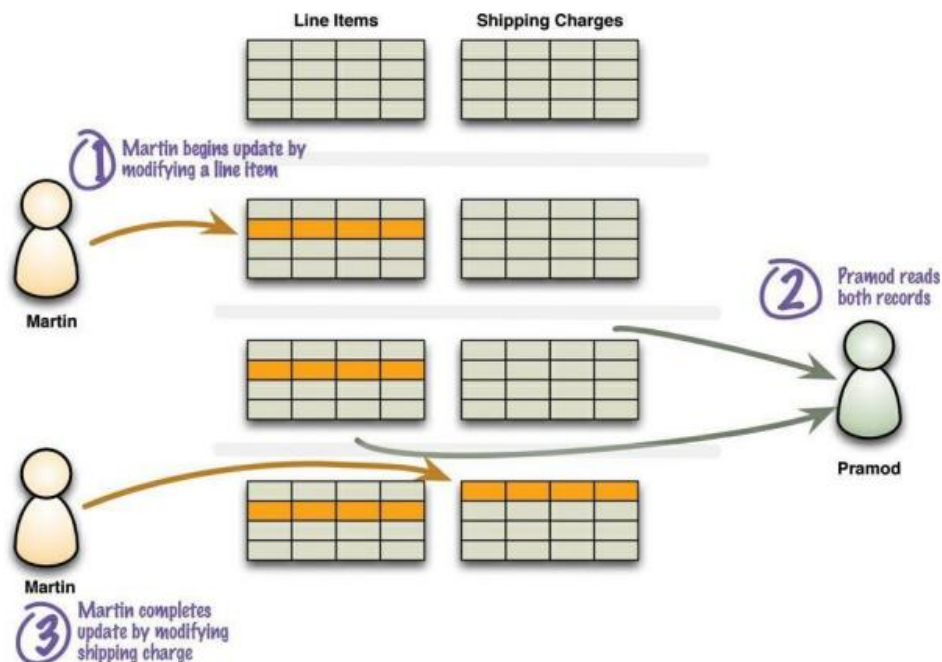


Figure 2.1. A read-write conflict in logical consistency

Q. 6 a) Explain briefly impedance mismatch, with a neat diagram.

Ans.

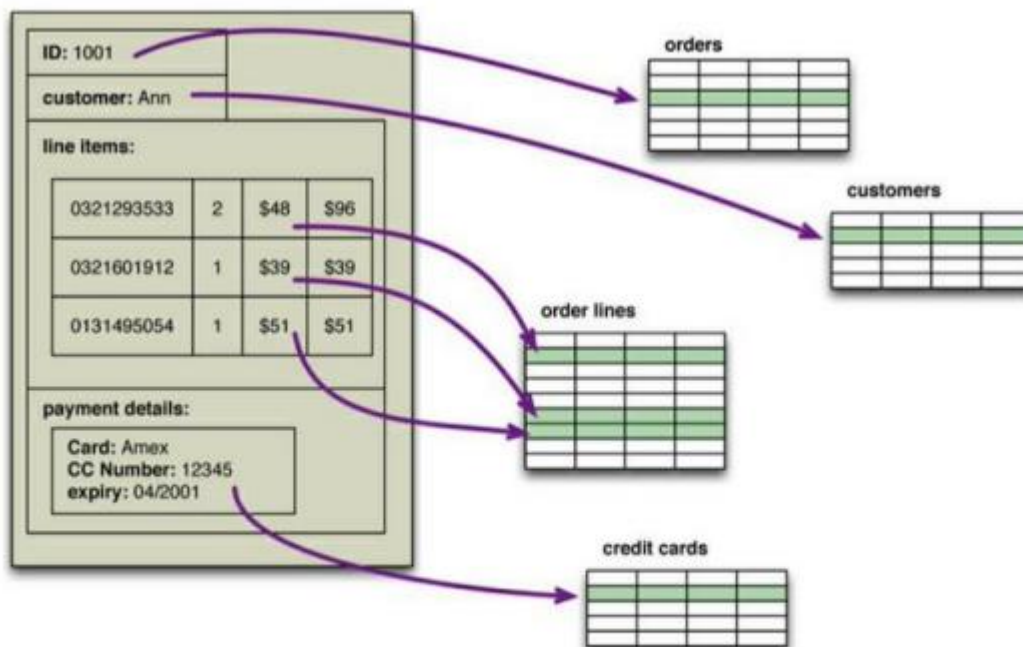
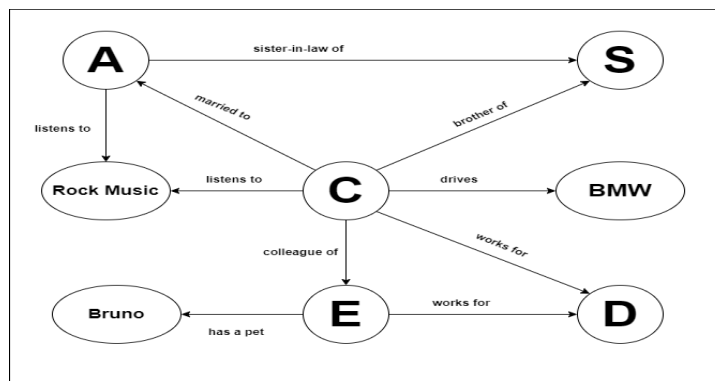


Figure 1: An order, which looks like a single aggregate structure in the UI, is split into many rows from many tables in a relational database



- For application developers, the biggest frustration has been what's commonly called the impedance mismatch: the difference between the relational model and the in-memory data structures.
- The relational data model organizes data into a structure of tables and rows, or more properly, relations and tuples.
- The values in a relational tuple have to be simple—they cannot contain any structure, such as a nested record or a list. This limitation isn't true for in-memory data structures, which can take on much richer structures than relations.
- If you want to use a richer in-memory data structure, you have to translate it to a relational representation to store it on disk. Hence, the impedance mismatch—two different representations that require translation.
- The impedance mismatch led to relational databases being replaced with databases that replicate the in-memory data structures to disk. That decade was marked with the growth of object-oriented programming languages, and with them came object-oriented databases—both looking to be the dominant environment for software development in the new millennium. However, while object-oriented languages succeeded in becoming the major force in programming, object-oriented databases faded into obscurity.
- Impedance mismatch has been made much easier to deal with by the wide availability of object relational mapping frameworks, such as Hibernate and iBATIS that implement well-known mapping patterns, but the mapping problem is still an issue.
- Relational databases continued to dominate the enterprise computing world in the 2000s, but during that decade, cracks began to open in their dominance.

6 b) Use the above diagram and answer the following questions.



- Who listens to rock music and works for D?
- Who works for D and has married to colleagues?
- Who listen to rock music?
- How are A and S related to each other and also to C?
- What are the genders of C and A?

**Ans.**

a) C

b) E

c) C, A

d) A is married C and sister-in law to S

e) A is female & C is Male.