BCS Higher Education Qualification

Diploma

April 2023

EXAMINERS' REPORT

Database Systems

Questions Report:

A1	
	This was a popular question, with around half of candidates achieving a pass mark. Part a) was fairly well answered, though few candidates provided sufficient detail of how data independence is achieved when the physical layer is modified and existing applications remain unaffected. There was also a lack of examples of some of the changes that can be made without affecting existing applications. At least four examples were expected. One of these could be, for example, modifying physical file organisation. Part b) concerned related fundamental relational database concepts which were fairly well understood but often lacked clarity and sufficient detail in explaining the differences. Many candidates did not provide any answer to subpart iii). The strongest answers were from those candidates that provided clear examples when answering compare and contrast questions.
A2	
	This was one of the least popular questions with around half of candidates making an attempt. However, more than half of candidates gained a pass mark. Part a) was generally well answered though some candidates did not address the context of the question. This concerned maintaining consistency by ensuring accuracy, correctness and validity given various threats that can cause a loss of integrity. Part b) included a different type of integrity concerning domain integrity. Most candidates showed a good understanding that a database needs to follow defined rules for values, range, and format. A description (by use of examples) of how these rules are enforced to validate, or restrict data was expected in order to gain full marks. Part c). This part was generally well answered with most candidates understanding the concept of recovery and concurrency control of interleaved transactions. Most of the marks were awarded for a timeline diagram consisting of three or four transactions that run concurrently. Candidates were generally aware that a checkpoint marks the start of the series of transactions and records this in the transaction log. The timeline diagram also marked the point of failure and from that an explanation of how recovery can be performed by going back to the last checkpoint before failure and restoring/ running transactions that were unsafe from the transaction log. Part d), this question produced a range of answers and proved to be a good test of applying knowledge to a realistic catastrophic event. Candidates found it challenging

	to produce the stages required when assessing the risks of complete data loss. Very few candidates could identify all the risks and write a series of actions with explanations. Most candidates could recommend sensible solutions that protect or foresee such events, such as distributing data resources (for example replication), off-site backups (a backup site in a different geographical region) and Cloud Computing (outsourcing).
A3	
	This was not a popular question with a third of candidates making an attempt. However, candidates performed well. This question was seemingly chosen by candidates fluent or confident in SQL. Candidates with good SQL knowledge, mainly through practice, are able to gain good marks on this type of question. Part a) was generally well answered with many candidates gaining full marks. A good understanding of how SQL keywords (in particular SUM, GROUP BY, ORDER BY INNER JOIN) are applied to the sample query. RIGHT/LEFT/INNER JOIN were generally well understood and the differences explained effectively using Venn diagrams. Part b), approximately half of the candidates succeeding in gaining high marks. It is recommended to show the mechanism of a more complex SQL query as this can attract marks even if the solution is incorrect.
B4	
	Three quarters of students attempted this question, making it a fairly popular choice, however less than half of the attempts achieved a pass grade. Part a) was usually answered, but often without precision: 'BCNF is a NF that is not often used' is not a sufficient answer, slightly better answers identified BCNF as 3.5 NF or stated that it comes after 3rd NF – but a good answer would have needed to explain what it adds (i.e. the removal of non-trivial functional dependencies). Part b) - In the past, attempts were often received showing 1st NF, (which needs only to remove repeating groups), already designed to be a number of linked tables. This was less common this session, but still occurred in some answers. The designs of second and third NF however did not always identify elements that were functionally dependent on each other when deciding what to retain in one table – many answers included a single table with movies and addresses, while separating people's names into a table of their own. Many candidates did not attempt part c). The answer required a small number of SQL DDL statements (create table) for the design created earlier. Where candidates answered, they did well by producing good SQL statements. The most common errors included missing the definition of keys or having the foreign key in the wrong table (i.e. in such a way that it would require multiple PK values to be inserted in each FK field to allow pointing to all linked rows in the other table).
B5	
	This question was answered by a large majority of the candidates. Over half of candidates achieved a pass grade in the question. Part a) required a few definitions of database concepts with examples for each. Most candidates provided clear definitions for attributes, tuples and one-to-many relationships giving suitable examples. Definitions for functional dependencies and modalities often did not convey a clear understanding of the concepts and also often

	lacked an example. Especially for functional dependencies, explanations often spoke about primary keys, but the primary key might just be one of a set of candidate keys so the dependency concept is wider. Part b) required candidates to explain a number of diagrammatic notations; answers often were not given or were imprecise by not explaining the detail of 0 (zero) or many or one to one and how they would be used and what they would mean in the design. Parts c) and d) addressed relation algebra and joins, with some questions left unanswered by many students (especially part d). Those who answered, generally did well on translating the first two algebra statements to an SQL query, but struggled with the third along with the query in part d) by selecting the wrong type of join statement.
B6	
	This question was answered by a majority of the candidates. Candidates generally provided good answers leading to higher marks. Part a) required the design of an ERD (or other suitable) diagram for a database design; almost all candidates identified the right entities and the respective attributes from the description. Genre was not always identified as a separate entity. Part b) required the design of a set of tables to match the diagram, and again answers were generally correct with some minor errors. A common mistake was the embedding of foreign keys in the wrong table. For example, the 1:n relationship between *Game* and *Genre.* Each game has a genre, but there can be many games of the same genre. This requires the foreign key in the *Games* table pointing to the *Genre* table. Some candidates did not mention keys in the design.