BCS Higher Education Qualification

Diploma

October 2023

Examiners' Report

Database Systems

Questions Report:

Marker name	All markers of section A are required to add comments here:
A1	
	This question was attempted by only 30% of candidates, of which only 20% passed. This question was poorly answered, with candidates seeming to know that not all databases are relational, but with very limited understanding of the systems beyond the relational sphere. Parts a, b c and d) highlighted this very clearly with answers being superficial, incomplete (some parts not answered) or simply presenting ideas outside the field of the question. Part e) presented a lack of understanding of big data and the key aspects that underpin it.
A2	
	This question was attempted by almost all candidates, with 2/3s achieving a pass grade in the question. Candidates scored good marks on this question, which focused on normalisation. Many candidates realised that the table was already in 1NF, there was good understanding of the theoretical ideas behind each NF but many candidates did arrive at schemas that were between 2nd and 3rd instead of only resolving to 2nd NF before moving to 3rd. Some candidates did not present interim steps and attempted to present a 3NF table set straight way, which often went wrong in some detail. Definitions of BCNF were often vague and the understanding of functional dependencies was often not explained in sufficient clarity to reassure that it was understood.
A3	
	This question was attempted by 2/3s of candidates, with about half of those achieving a pass grade. Answers in this question were often correct on the factual base of the result of queries, but the accompanying explanations were not good. Were candidates had to explain the purpose of queries they resorted to explaining each line in its own rather than seeing the overall purpose. While it is important to understand the code, it is also very important to be able to express the purpose of queries at the business level.

Marker name	All markers of section B are required to add comments here:
B4	
	Part a) A fairly popular question attempted by around 78% of candidates. Overall performance was slightly below par, being just below the pass rate of around 50%. There was a good spread or marks. Candidates needed to reference the comic/book scenario in their answers and those candidates who did this successfully gained the highest marks. More time needs to be spent reading/understanding a given scenario and understanding the context because it provides in examples in all the parts of this question. This part compared the different approaches to data representation and manipulation. Many candidates' answers lacked depth and had difficulty expressing the key features often omitting important detail such as data redundancy and consistency; visual appeal; ease of presenting data; presenting images. An obvious comparison is that a spreadsheet presents data exactly how it is stored whereas a database approach separates the logical and physical view (tables) from the user view (as a spreadsheet). Data manipulation. Most candidates covered the significance of having SQL to query and manipulate data in comparison to a spreadsheet that used built in search and functions to the spreadsheet. Clearly the scenario had a need for simplicity in adding and manipulating data which a spreadsheet approach provided. The ability to access different data sets (called sheets in Excel) would satisfy the needs of the comic book narrative. To emphasise; the best answers addressed the features to compare in the context of the classic comic scenario.
	Part b) Generally covered fairly well but again answers lacked detailed examples related to the scenario. Some candidates produced very generic answers but most agreed that a file-based approach was a better approach given the support that a spreadsheet approach applies.
	Part c) The scenario was extended to include trading of comics to customers. The ability to scale up using a spreadsheet approach is more problematic and most candidates covered this. However many answers lacked an explanation of why a database approach is better at scaling up than a file-based approach. Validation and data integrity and Authentication issues were generally covered very well.
	Part d) Mostly generic answers as candidates could recall the essential features of data independence. A number of candidates provided examples that were already covered and pertinent to the question parts above.

B5	
	This was the most popular question on the paper attempted by around 90% of candidates. Overall performance was the best on the paper with a pass rate of around 75%. There was a good spread or marks.
	Part a) Most candidates understood the different ER modelling symbols and what they conveyed. The Crow's foot notation expressed one (mandatory) and one (mandatory) that is one and only one just in case there is any confusion.
	Part b) Most candidates answered this part well although marks were lost as follows. This part required an understanding of what is meant by a physical data model and the process of mapping - how a logical (ER) model is translated to a physical design. Some candidates lost marks if they used the name attribute as identifiers/Primary keys of the respective tables rather than introduce a primary key such as EmpId and DepID. Many candidates lost marks because they did not specify the primary keys and identify the foreign keys and connections between the tables.
	Part c) Generally good knowledge of SQL DDL statements was evident, particularly pleased to see syntactically accurate code which was not absolutely essential.
	Part d) Again well answered most candidates were familiar with adding check constraints into SQL DDL code.
B6	
	An unpopular question, attempted by around 35% of candidates. Overall performance was slightly below par with a pass rate of around 44%. There was however a good spread or marks with a number of candidates achieving 25 marks.
	Part a) Most candidates identified that the interleaved transactions caused the write from transaction TX1 being lost. Many candidates appreciated the need for using a locking strategy to avoid a lost update. But most candidates revealed a limited understanding of applying 2 phase locking to the transaction schedule.
	Part b) As expected most candidates were familiar with the key requirement of backing up and gained high marks.
	Part c) Auditing was less familiar to candidates particularly how an audit trail can track, and record changes made to a database.
	Part d) Generally, candidates found this part difficult with some candidates not providing any answer. The SQL keywords GRANT and REVOKE required an accurate interpretation of the access rules to three different users Sam, Jane and Joe. To ensure consistency all permissions were revoked before granting access to each individual. Part d (iii) was attempted successfully only by a small number of candidates.