Examiner Report	
Qualification Name	Higher Education Qualification
Qualification Level	Diploma
Date/ Series	April 2024
Module	Database Systems
Question no.	**comments**
A1	Syllabus: 1.1 (describe theoretical concepts) This question was very popular with 85% of candidates addressing it of which 57% achieved a pass grade. Overall answers were of reasonable standard, there was a fair number of candidates not addressing some of the sub-questions. Part a) considered definition and examples for basic terminology; most candidates could identify the concepts correctly but often an example was missing. The concepts of composite keys and schema were not well explained by many. In part b) the discussion on data independence support provided by the ANSI SPARC architecture sometimes lacked the focus on data independence (students either discussed database design or SPARC but did not really answer the question); some questions were just a collection of facts without a clear relation to which are relevant to the answer of the question. In part c) the question asked for a number of relational operators, generally it was answered well but often the answer only included a listing of the operators without application to the sample tables given.
Question no.	**comments**
A2	Syllabus: 4.1 (normalisation) This question was answered by 75% of candidates of which 70% passed. Part a) asked for issues that can affect unnormalised tables; these were generally well identified and good examples given. Candidates scoring low on this question did not answer it, but rather provided an outline of how normalisation would be undertaken.

	In part b) candidates generally did well in normalising the tables, but they often continue to struggle with the concept of a 1NF with frequently providing a normalisation to 2NF instead which then, as a consequence, leads to confusion as to further normalisation steps. In part c) most candidates were able to identify a clear solution to adding the additional requirement. Part d) Identifying dependencies was often not answered; where it was, the dependencies frequently identified the right elements but not the right dependency relation.
Question no.	**comments**
A3	Syllabus 3.2 (relationship constraints), 5.1 (SQL) 43% of candidates answered this question, of which 80% passed. So overall candidates did very well in this question. Part a) considered entity integrity and its formulation in SQL – most candidates could explain the required concept and provided correct SQL for it. Part b) considered referential integrity (again concept and SQL) and candidates provided robust answers; there were slightly more candidates struggling with explaining the necessity of referential integrity. Part c) The concept of a NOT NULL constraint was identified by some, but not all candidates. Part d) required the creation of a SELECT query, which candidates generally managed well but some found the idea that there were two nested queries to create a difference set challenging. Part e) required discussion of advantages of text and GUI interfaces; this generally led to the weakest answers in question A3 with many students focusing on colour schemes rather than the user skill required to use on or the other and the functionality readily accessible through one vs the other.

Question no.	comments
B4	Syllabus 2.1, 2.2 This was a fairly popular question attempted by 42% of candidates with an average mark with around 40% gaining a pass mark. Overall many candidates lost marks because they gave very superficial answers with little reference to the scenario to draw examples (apart from part d). Candidates are advised to thoroughly read through the scenario before answering this type of question. Part a) Some of the weaker attempts were from candidates that only provided generic background knowledge. In addition candidates needed to provide relevant examples related to the scenario in order to gain the highest marks. However it was apparent that many candidates were unfamiliar with Graph databases and therefore could not expand or apply the use of Graph databases to the scenario. For example route planning and tracking were efficient and effective handling of complex relationships was required. The link to NoSQL databases was often mentioned in passing and carried over or simply repeated in part c) Part b) Many candidates focused on generic knowledge comparing relational databases versus flat files without reference to the scenario. Generic answers included limitations of flat files in comparison to Relational databases. The best answers were from candidates who provided a balanced view discussing the pros and cons of document oriented storage which can be a more flexible way of storing timetable data. For example considering the various downsides of handling semi-structured storage in comparison to structured data as per relational databases. The benefits of a document oriented database, such as maintaining the natural flat file representation of timetable data could outweigh the limitations that were generally well understood by many candidates. Part c) This part expanded part a) requiring discussion on the merits and drawbacks of two contrasting approaches. Again there was a lot of repetition of points already made and again some lack of knowledge of graph databases what they do and why they might be applied to the scenario. Overall though only a small number of answers were able to gain high marks. The interpretation of the scenario varied amongst candidates to such an extent that no definitive answer to the question was forthcoming. Most important was to address the key reasons why

	one approach was better than the other and in the main were not adequately addressed. The key points might include balancing the need for structured data to accommodate all round support from a RDMS to provide efficient querying, updates versus the flexibility and ease of naturally modelling graph/map data that is very difficult in a relational model. Part d) This part was answered fairly well overall with standard text book responses that covered the main points. Candidates generally gained high marks on this part.
Question no.	**comments**
B5	Syllabus 3.1 3.2, 5.1 This was a very popular question with over 95% of attempts. Most candidates exhibited little difficulty in gaining high marks with a pass rate of almost 80%. Part a) Generally well answered. Many candidates provided examples of the various notations to good effect. In particular examples often avoided the confusion that some candidates had regarding the concepts of cardinality and participation constraints in relationships. Most candidates identified that “total cost” was a derived attribute not a standard attribute such as “Price”. Part b) This part was generally well answered with most candidates identifying the primary and foreign keys. A significant number of candidates omitted the composite keys required for the Enrolment table. Many candidates simply rewrote CREATE TABLE statement with the addition of keys, rather than write ALTER TABLE statements This was much better practice and was rewarded as a result. Part c) Almost all candidates noticed the Tables in part b) represented the same Entity Types Course and Student. All they needed to do was to resolve the many to many relationship to reflect which represented the Enrolment Table in part b) as a Table called “Takes”. Despite this clear steer, a small number of candidates showed they were unfamiliar with mapping ER models to Tables. Part d) The NOT NULL clause in SQL DDL was well understood by most candidates and suitable examples applied to an attribute/column were provided. It should be noted that most DBMS imply NOT NULL when a Primary key is defined in a DDL CREATE TABLE statement. Therefore part of the DDL in part b) could be rewritten as:

	CREATE TABLE Student (Name VARCHAR(50), DoB DATE NOT NULL, Studid INT, PRIMARY KEY (StudID))); Note that DoB must accept Date values whereas Name does not.
Question no.	**comments**
B6	Syllabus 6.1, 6.2 This question was a fairly popular question attempted by around 57% of candidates with a similar amount of candidate passes. Part a) Most candidates correctly identified the problem caused a result of interleaving two transactions without any concurrency control. A Dirty Read or Uncommitted Update had occurred because transaction TX2 read and wrote a value for "x" which had been subsequently roll backed. 2 Phase Locking was familiar to most candidates as they could describe the protocol fairly well. About half of candidates were able to describe how a Dirty Read described in the scenario could be avoided. Part b) A variety of answers were received with some candidates misinterpreting what was required. A typical response from candidates was to only look at standard recovery measures such as local backups rather than consider a strategy for disaster recovery; if say local backups are destroyed by fire or theft for example. Some candidates focused on cyber attacks such as hacking again a valid plan would be necessary to support a business to quickly regain access to its data. Typical approaches include duplicate data centres; outsourcing to the cloud; the ability to transfer operations to a remote site. The strategy should identify which systems are the most critical and hence in which order they need to be recovered to minimise outage/disruption to the business. Part c) Overall most candidates indicated that an alarm should be raised if there has been attempts to change login credentials, possibly to gain access by an unauthorised user. Most candidates expressed the serious nature of a possible cyber attack in the worst case when changes in login/access permission had occurred. Some candidates went further and gained the highest mark by fully addressing what the initial reaction to such an attack prior to responding/remedial action needed. For instance, there could have been a cyber attack on the DB server or a hacking attempt to

	gain access rights. Also there may be an employee that had access rights that should not have been granted. Part d) This part had three subparts. Most candidates answered the first two subparts very well. Many candidates realised that the use of roles avoids the need to assign permissions to individual users as candidates should have noted in their answer to subpart i). The third subpart required knowledge of using Views in the context of access control and security. The main points that were often overlooked in candidate answers were that access rights can be granted on views in the same way that they can on tables; however a view allows an extra level of abstraction and hence can further hide information in for example specific columns that are not relevant to the given user/role. Views can combine information from various tables reducing the number of objects that a user/role requires access to. They also provide a primitive form of row level security.