Question number: A1
Total marks allocated: 25
Examiners' Guidance Notes
This question was attempted by a majority of candidates, but the popularity did not match the overall performance which was disappointing with an average mark of 8.4. The weak performance may be attributed to a perception that recalling standard text book knowledge was sufficient. Many candidates produced rather shallow knowledge often not relevant to the question topic. Part a) required a balanced explanation that focused on the different nature of file storage in operating systems (either Linux or Windows) with database systems. Instead many candidates reiterated standard bookwork on advantages of a database system over a file based system Part b) revealed a similar pattern with a general lack of depth in the answers. For example ANSI standards was often focused on the single topic of ANSI SPARC architecture and only a small number of candidates seemed to be aware that ANSI standards are applied to other areas such as the SQL standard. Knowledge of schemas was especially poor with quite a lot of confusion in the answers.

Question number: A2
Total marks allocated: 25
Examiners' Guidance Notes
This question was quite popular, attempted by around two-thirds of candidates. The performance was reasonably good with just under half of candidates obtaining a pass mark. Some candidates produced answers on database integrity that covered a wide range of topics that did not reference data protection at all – the overall context of part b). Answers to Part c) were generally less successful with many candidates unfamiliar with modern and more automated approaches to database administration. But with 3 marks available this made a small impact on overall performance.

Question number: A3
Total marks allocated: 25
Examiners' Guidance Notes

This was a popular question attempted by around half of candidates. The average mark of 11.4 indicates a good level of performance with around 60% gaining a pass mark.

Part a) was generally less well answered with numerous inaccuracies including some needless arithmetic errors when calculating the formula.

Part b) candidates appear to be familiar with this topic concerned with Relational Algebra (RA) and could therefore readily recall RA operations. Most candidates elected to draw Venn diagrams (the most natural way to illustrate RA operations). The application of RA operations to the data supplied was required to gain full marks though.

Part c) again was generally well answered but there were some candidates who misread the question and produced too much irrelevant detail on every type of Join operation they could conceive.

Question number: B4
Total marks allocated: 25
Examiners' Guidance Notes

This question was very popular with the vast majority of candidates attempting it. The average mark was around 50% for the question. Candidates engaged well with creating SQL statements for creating the two tables, but struggled somewhat with defining the foreign key. Candidates generally did well on the normalisation process, arriving at a variant that was close to being in third normal form. However many candidates did not produce a first normal form, but went straight to 2nd calling it 1st. Also, several candidates described only what a normal form looks like rather than applying it to the example. The answer to the usefulness of denormalization was disappointing with most candidates seeming to be under the impression that this is to be done to make the database more readable or to save on storage space.

Question number: B5
Total marks allocated: 25
Examiners' Guidance Notes

This question was fairly popular with candidates with 62.2% attempting it.

Most candidates managed well to define the 1:n relationships, but many failed to construct a diagram with an intermediate table for the n:m relationship.

The definition questions for the 5 terms seemed to be surprisingly quite difficult for candidates; with very few being able to explain several of them. The concept of cardinality ratio and participation constraint seemed particularly troublesome.

Question number: B6
Total marks allocated: 25
Examiners' Guidance Notes

This question was fairly popular with candidates with around two thirds of candidates attempting it, gaining on average about 50% of the marks.

Most candidates did well in identifying the main tables from the description and linking these together. However, cardinalities were often not captured well (they sometimes were not shown at all, or shown the wrong way round) and any n:m relationships caused trouble for most. While there was a choice of notation, it was sometimes used inconsistently (that is,.a mix of UML and the crowsfoot notation). When translating the diagram to tables, the keys were quite often in the wrong table (for example, a key to trailer in the trailer park table, rather than the other way round), foreign keys were forgotten, keys occurred in two tables cross referencing each other; fields were sometimes included in the wrong entity (for example, rent was held within the manager table).