

Review

CSC230 Database Technologies for Analytics

09 November 2016

1. In his Turing Award lecture, Michael Stonebraker explained that he wanted to create a database management system that could store more than words and numbers. What other kind of data did he think was important?
2. The largest value that can be stored in a field whose type **INTEGER** is 2,147,483,647. Why is the upper bound not an easier to remember value like 999,999,999 or 1,000,000,000 or 1,000,000,000,000?
3. Computability, complexity, and correctness are three important concepts and subjects of study in computer science.

Which of these three “C” words best applies in each of the following cases?

- (a) With some sorting algorithms, a computer will execute a number of instructions that is roughly proportional to N^2 , where N is the number of items in the list that is being sorted. With other sorting algorithms, a computer will execute a number of instructions that is roughly proportional to $N \log N$.
 - (b) There are computer programs that control vehicles, power plants, and robotic surgical instruments. The software engineers who write these programs bear responsibility for the protection of lives and property.
 - (c) Alan Turing proved the impossibility of writing a program that can check any other program to determine whether or not its execution with a given input will terminate in finite time.
4. What are the advantages of using a file system as a database?
 5. What are the disadvantages of using a file system as a database?
 6. An ACID database management system has four properties:
 - Atomicity
 - Consistency
 - Isolation
 - Durability

Which of these four properties applies in each of the following cases?

- (a) The system can retrieve/recover data in any circumstances, even after the failure of hardware.
 - (b) The system does not allow to one operation to see the partial (incomplete) results of another operation.
 - (c) The system does not allow any operation that does not respect constraints that define the integrity of the database. For example, it forbids the deletion of a row in one table when a row in another table contains a reference to the first row.
 - (d) The system erases/reverses all effects of any operation that it is unable to complete.
7. What is “brittle” software?
8. What are some advantages of using a declarative language?
9. What, according to Philip Greenspun, does a caveman programmer need to know about relational databases?
10. What is the value of an index in an RDBMS?
11. (a) Give an example of a SQL statement that is part of the Data Definition Language?
- (b) Give an example of a SQL statement that is part of the Data Manipulation Language?
12. Philip Greenspun identifies an alternative to relationship databases and a drawback to that alternative.
- (a) What is the alternative?
 - (b) What is the drawback?