一 選擇題（每題 3 分）

1) Which of the following is **not** a reason to explain why the ER model has been so successful in conceptual DB design?
   A.) It is simple.
   B.) It gives us a very natural way to organize the information.
   C.) It integrates well with other DB concepts.
   D.) It has an intuitive appeal.

2) The data in a DBMS is often described at three levels of abstraction, namely external schema, conceptual schema, and physical schema. Such an abstraction gives us an important advantage of using a DBMS, which is
   A.) Execution efficiency
   B.) Data independence
   C.) Improved data integrity
   D.) Reduced data redundancy

3) Which of the following formal query languages is operational in natural?
   A.) Relational calculus
   B.) Relational geometry
   C.) Relational algebra
   D.) Relational topology

4) Which of the following is a special case of the join operation on two relations R and S in which equalities are specified on all fields having the same name in R and S?
   A.) Outer join
   B.) Semijoin
   C.) Natual join
   D.) None of the above

5) In database recovery, a common property called write-ahead log (WAL), is often maintained. What does it mean?
   A.) The on disk log must be written with a special value before its use.
   B.) Each write action must be recorded in the log on disk before the corresponding change is reflected in the database itself.
   C.) One must create a log on disk before we can use the recovery function of the DBMS.

(背面仍有題目,請繼續作答)

D.) None of the above.

6) XML has the following features except
   A.) It allows users to define new collections of tags.
   B.) It has the potential to make database systems more tightly integrated into Web applications than ever before.
   C.) It was developed to have much of the power of SGML while remaining relatively simple.
   D.) It replaces HTML.

7) Which of the following is not a difference between main memory and disks?
   A.) Main memory accesses are faster than disks.
   B.) The time to access a disk page is constant.
   C.) Data in a disk remain intact after a power failure.
   D.) Each byte of main memory is more expensive than that of a disk.

二 問答題

1. Define the following terms: functional dependency, superkey, referential integrity, relational completeness. (12 分)
2. What is the difference between a candidate key and the primary key for a given relation? (3 分)
3. Give a set of functional dependencies for the relation schema R(A,B,C,D) with primary key AB under which R is in 2NF but not in 3NF. (3 分)
4. DB transactions often have the so-called ACID properties: atomicity, consistency, isolation, and durability. Please explain each of the four properties. (12 分)
5. A company has many employees and departments. Every employee is identified by his/her social security number (ssn), and has a name and an address. Every department is identified a department id (did), and has a name and budget. Many employees could work in a department. An employee could work in many departments. Answer the following questions.
   (1) Draw an ER diagram to express the scenario described above. (5 分)
   (2) Write SQL statements to create the corresponding relations for the ER diagram you draw (choose appropriate type for each attribute in your relations) and capture as many of the constraints as possible. (9 分)
   (3) Write an SQL assertion to express the following total participation constraint: Every department must have at least one employee. (4 分)

6. Consider relation R with five attributes ABCDE. You are given the following dependencies: A → B, BC → E, and ED → A.
   (1) List all keys for R. (6 分)
   (2) Is R in 3NF. Why? (3 分)
   (3) Is R in BCNF? Why? (3 分)
7. Consider the following schema:
   Suppliers(sid: integer, sname: string, address: string)
   Parts(pid: integer, pname: string, color: string)
   Catalog(sid: integer, pid: integer, cost: real)

   The key fields are underlined, and the domain of each field is listed after the field name. The Catalog relation lists the prices charged for parts by Suppliers. State what each of the following queries compute: (π denotes projection, σ denotes selection, |><| denotes equijoin, and ∩ denotes intersection) (9 分)

   (1) $\pi_{sname}(\pi_{pid}(\sigma_{color='red'} Parts) |><| (\sigma_{cost<100} Catalog) |><| Suppliers)$
   (2) $\pi_{sname}(\pi_{sid}(\sigma_{color='red'} Parts) |><| (\sigma_{cost<100} Catalog) |><| Suppliers)$
   (3) $(\pi_{sname}((\sigma_{color='red'} Parts) |><| (\sigma_{cost<100} Catalog) |><| Suppliers)) \cap (\pi_{sname}((\sigma_{color='green'} Parts) |><| (\sigma_{cost<100} Catalog) |><| Suppliers))$
8. Consider the following relations containing airline flight information:
   Flights(flno:integer, from: string, to: string, distance: integer, departs: time, arrives: time)
   Aircraft(aid: integer, aname: string, cruisingname: integer)
   Certified(eid: integer, aid: integer)
   Employees(eid: integer, ename: string, salary: integer)

   Note that the Employees relation describes pilots and other kinds of employees as well; every pilot is certified for some aircraft (otherwise, he or she would not qualify as a pilot), and only pilots are certified to fly. Write the following queries in relational algebra and SQL. If you can not express any of the following queries, explain your reason.

   (1) Find the total amount paid to employees as salaries. (2+3 分)
   (2) Find the eids of pilots certified for some Boing aircraft. (2+3 分)