1. (14 marks) For each of the following, circle the ONE, BEST answer.

   (a) Which of the following can be part of the database state?
   i) ER-diagram
   ==> ii) table
   iii) index
   iv) disc

   (b) Which of the following can be part of the internal schema?
   i) ER-diagram
   ii) table
   ==> iii) index
   iv) disc

   (c) Physical Data Independence refers to:
   i) the ability to add more database packages to the system without affecting applications
   ii) the ability to add more data models to the database package without affecting applications
   iii) the ability to add more tables to the database without affecting applications
   ==> iv) the ability to add more indexes without affecting applications
   v) the ability to add more discs to the computer without affecting applications

   (d) In the relational model, another word for relation is:
   ==> i) table
   ii) row
   iii) column
   iv) extent
   v) both a and d

   (e) In the relational model, another word for attribute is
   i) table
   ii) row
   ==> iii) column
   iv) extent
   v) none of the above

   (f) Concerning a B^+-tree, which of the following is **false**?
   i) it can be used to look up a record based on its key value
   ii) it can be used to build an index for a unique attribute
   iii) it can be used to build an index for a non-unique attribute
   ==> iv) it can be used to cluster records based on some hash value
   v) its structure tends to yield short, wide trees

This ER-diagram is to be used for the following 4 questions.

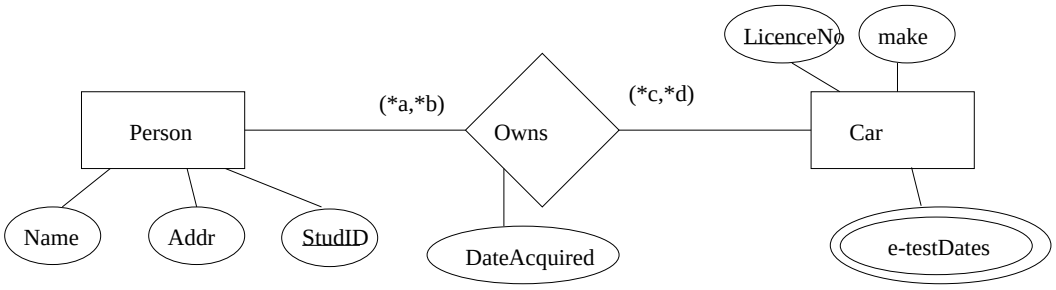

(g) "e-test dates" in this ER diagram is an example of:
i) an entity
ii) a relationship
iii) a simple attribute
iv) a compound attribute
==> v) a multi-valued attribute

(h) The correct substitution for (*a,*b) in the above ER Diagram is:
i) (0, 1)
ii) (1, 0)
iii) (1, 1)
iv) (1, n)
==> v) (0, n)

(i) The correct substitution for (*c, *d) in the above ER Diagram is:
i) (0, 1)
ii) (1, 0)
==> iii) (1, 1)
iv) (1, n)
v) (0, n)

(j) The insurance company wants to record who regularly drives each car. This information should be added to the ER Diagram as:
i) another attribute of Person
ii) another attribute of Car
iii) another attribute of Owns
iv) another entity for Driver
==> v) another relationship between Person and Car

(k) A foreign key in relation R
i) uniquely identifies all the tuples in relation R
ii) refers to the primary key of relation R
==> iii) uniquely identifies all the tuples in some other relation S
iv) refers to the primary key of some other relation S

(l) Consider the following ER diagram.

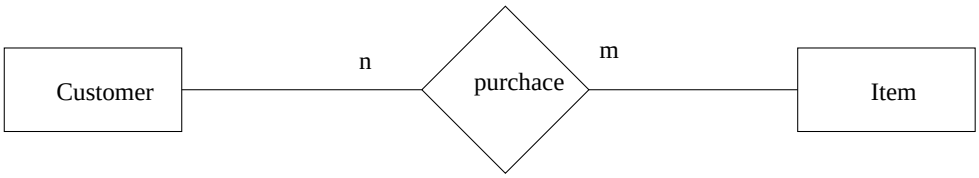

If we wish to add an attribute Date to this ER diagram, where should it go?
i) as an attribute of Student
ii) as an attribute of Desk
==> iii) as an attribute of assign

(m) Suppose we have data records which are 128 bytes each, disc pointers are 4 bytes each and disc blocks (pages) are 512 bytes each. When using hashing on discs, buckets are filled with records and then linked to overflow locations. How many records can fit in each bucket for a hash table?

i) 2
==> ii) 3
iii) 4
iv) 5
v) none of the above

(n) Suppose our records for a B^+-tree are 128 bytes each, the pointers are 4 bytes each and the keys are 4 bytes each. The disc blocs (pages) are 512 bytes. How may KEYS will fit on each internal page of this B^+-tree?

==> i) 63
ii) 64
iii) 65
iv) 3
v) 4
vi) none of the above

2. (7 marks) Consider the following B$^+$-Tree:

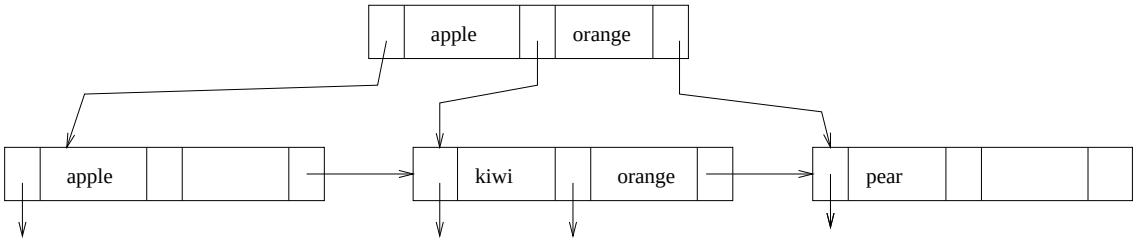

(a) Show the tree that results after inserting "acai".

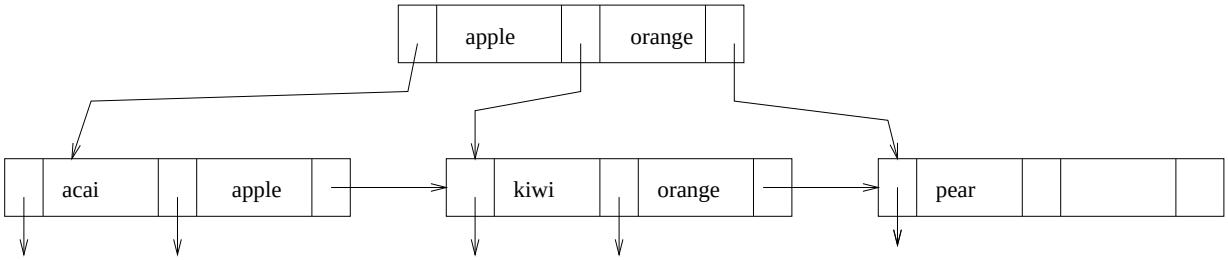

1 mark for leaf being changed, including the new down arrow
1 mark for leaving the rest of the tree unchanged.

(b) Show the tree after inserting "melon" in the **original** tree above.

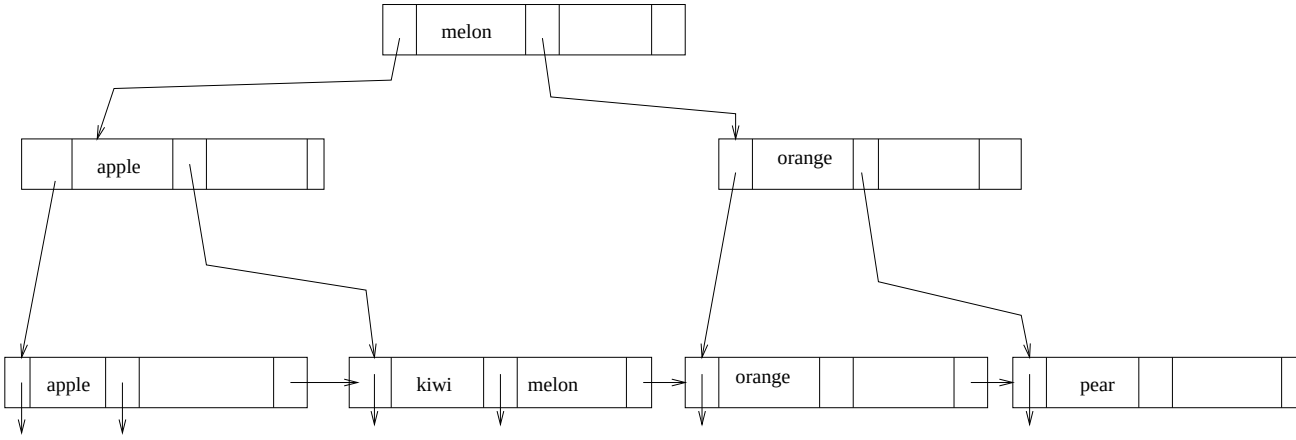

1 mark for new root node, 2 for correctly splitting the old root, 2 for correctly splitting the original leaf, and putting in all the down and horizontal arrows.

3. Consider the following relations:

R	A	B	C	D	E
	a	z	1	m	1
	b	x	1	m	6
	a	y	2	p	5
	b	z	2	p	6
	a	y	3	n	3

S	E	G	H
	6	1	a
	5	3	b
	2	5	c

T	H	I	J
	a	x	2
	a	z	2
	d	w	4

Show the relation (actual data) resulting from each of the following relational algebra queries:

(a) (4 marks) Result $\leftarrow \sigma_{C=2 \text{ or } B=\text{“}z\text{”}}(\text{R} \bowtie \text{S})$

	A	B	C	D	E	G	H
	a	y	2	p	5	3	b
	b	z	2	p	6	1	a

1 mark for right attributes, 3 for right data

(b) (4 marks) Result $\leftarrow \pi_{E,G,I,J}$ (S $\supset\!\bowtie\!\subset$ T)

	E	G	I	J
	6	1	x	2
	5	3	-	-
	2	5	-	-
	-	-	w	4

1 mark for right attributes, 3 for right data

(c) (6 marks) Show Temp1, Temp2 and Answer: Temp1 $\leftarrow \pi_C(\sigma_{E=6}(\text{R}))$; Temp2 $\leftarrow \pi_{B,C}(\text{R})$; Answer $\leftarrow$ Temp2 $\div$ Temp1

Temp1	C
	2
	2

Temp1	B	C
	z	1
	x	1
	y	2
	z	2
	y	3

Temp2 $\div$ Temp1	B
	z

2 marks for each table above

4. Consider the following relations:

   Customer(Name, Address, Age, Occupation)

   Item(ManuName, ItemName, Price)

   Likes(Name, ManuName, ItemName, Rank)

   (a) (6 marks) Give a **relational algebra query** to find, for all items made by Sony, the names of all customers who like something whose price is greater than 1000. Put the customer name and the item name in the answer.

   $\pi_{name,\ ItemName}\ (\ \sigma_{Price>1000\ and\ ManuName='Sony'}(\text{Likes} \bowtie \text{Item})$

   2 marks for the join

   2 marks for the select

   2 marks for the answer having the correct attributes

   (b) (6 marks) Give a **relational algebra query** to find all manufacturer, item name combinations which are not liked by anyone.

   $\pi_{ManuName,\ ItemName}(\text{Item})$ - $\pi_{ManuName,\ ItemName}(\text{Likes})$

   2 marks for getting 2 columns from Item

   2 marks for getting 2 columns from Likes

   2 marks for subtraction

   (c) (6 marks) Give a relational **algebra query** to find all customer names such that the customer likes ALL the items liked by "George".

   $\pi_{Name,\ ManuName,\ ItemName}(\text{Likes})$
   $\div\ \pi_{ManuName,\ ItemName}(\text{Likes})$

   2 marks for getting 3 attributes from Likes

   2 marks for selecting George and projecting onto the two attributes

   2 marks for the division

5. (10 marks) Give relation definitions for the following ER diagram. Underline primary keys and indicate foreign keys with a double underline.

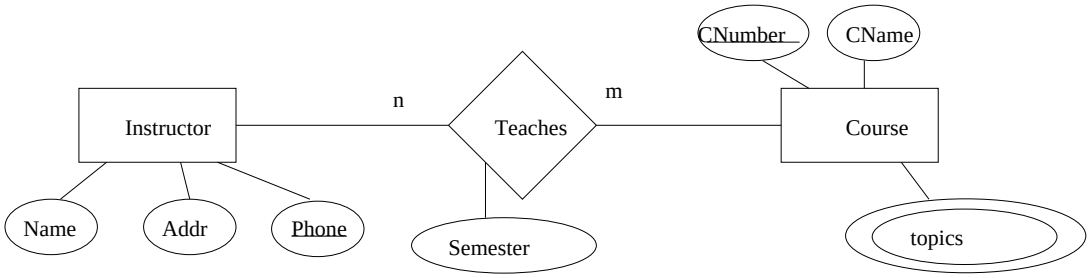

Course(CNumber, CName)

Instructor(Name, Addr, Phone)

Teaches(Cnumber Phone, Semester) or Teaches(Cnumber Phone, Semester)

Topics(CNumber, topic)

4 for the 4 tables (1 ea)

4 for the keys (1 ea)

2 for foreign keys