CSE-3421 Test #1

"Design"

Family Name: _____
Given Name: _____
Student#: _____
CS Account: _____

- **Instructor:** Parke Godfrey
- **Exam Duration:** 75 minutes
- **Term:** winter 2009

Answer the following questions to the best of your knowledge. The exam is closed-book and closed-notes. Questions will not be answered during the test. Should you feel a question needs an assumption to be able to answer it, write the assumptions you need along with your answer. If you need more room to write an answer, clearly indicate where you are continuing the answer.

There are four major questions worth 10 points each for 40 points in total.

Regrade Policy

- Regrading should only be requested in writing. Write what you would like to be reconsidered. Note, however, that an exam accepted for regrading will be reviewed and regraded in entirety (all questions).

Grading Box	
1.	/10
2.	/10
3.	/10
4.	/10
Total	/40

1. **E-R Design.** *Er, an E/R E-R, eh?* (10 pts) [exercise]

   You have been recently hired by Duey, Cheatem, & Howe, Inc., (DC&H) as a database designer. (Unfortunately, you have been assigned as an assistant to the infamous chief database designer Dr. Datta Bas. So your job will include correcting Dr. Bas's many logical mistakes.)

   DC&H is hired to construct a database for patient tracking in the emergency room of the Eiffel Badd Hospital. This has been assigned to your team. Specifications are as follows.

   An *in-patient* is someone who has been checked into the E/R (emergency room). The in-patient has associated a list of *traumas* (for instance, broken ribs, internal bleeding, and concussion). An in-patient may have one trauma or several. (These are the reasons why the in-patient is in the E/R!) The *severity* of each of the patient's traumas should be recorded.

   While in the E/R, various *procedures* will be performed on him or her. What procedures are performed on the in-patient should be recorded, along with the *success* level observed for each procedure in treating the in-patient's trauma and *when* the procedure was performed.

   Each procedure performed on the patient is conducted primarily by one *physician*. (Different procedures performed for the patient may be handled by different physicians.) Which physician handles each procedure should be recorded along with how much time (*ptime*) he or she spent (for billing purposes). Each procedure is associated with one of the in-patient's traumas, which justifies why it was performed. Of course (unfortunately), not every trauma of the patient may be treated by some procedure. A trauma may go untreated. Likewise, there can be *several* procedures performed on the in-patient for the same trauma. Furthermore, the same (type of) procedure might be performed several different times on the in-patient for the same trauma.

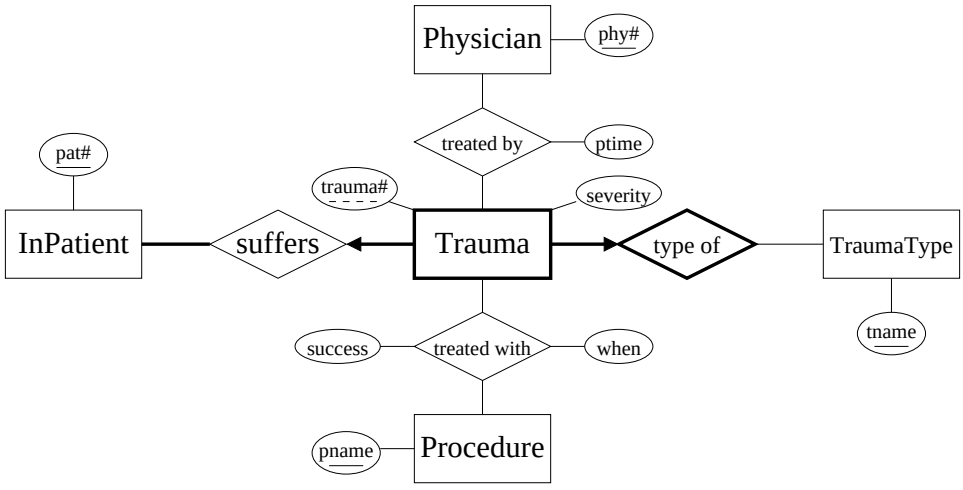

Figure 1: Dr. Bas's attempt of the E/R E-R.

a. (4 points) Figure 1 represents Dr. Bas's first attempt. Identify clearly at least two distinct major problems with the E-R diagram with respect to the specifications.

- *There is no way to know which physician did which procedure.*
- *The relationship* treated with *does not accommodate that the same procedure may be applied a second time to a given patient's trauma.*

b. (6 points) Present your own E-R design for the E/R specifications.

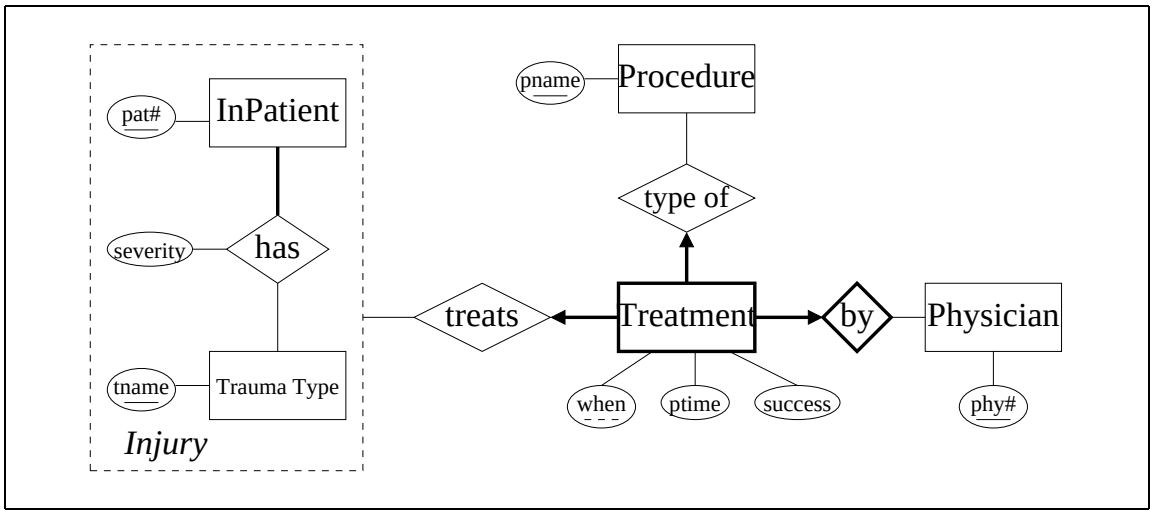

2. **General.** *I'm starting a general practise.* (10 pts) [True / False]

   Choose *T* (for *true* or *F* for *false* for each of the following. There is no deduction for wrong answers.

a. [F] It is impossible to represent a one-to-one relationship in a relational schema.

b. [F] Anything that can be expressed in an E-R diagram via a ternary relationship can be expressed in some other logically equivalent way without the use of ternary relationships.

c. [F] Under a relational database system, if table **R** has a foreign key constraint referencing table **S**, then each tuple in **R** is necessarily related to some tuple in **S** via the foreign key.

d. [T] The `SQL` statement "`DELETE FROM R`" might cause tuples in tables other than just **R** to be deleted.

e. [F] NULL values can be used to opt a tuple out of enforcement of the primary key.

f. [T] NULL values can be used to opt a tuple out of enforcement of a foreign key.

g. [F] One of the tasks of a relational database management system (RDBMS) is to ensure that relational schemas are in at least 3NF.

h. [T] Consider a relation with four attributes (say, A, B, C, and D). There are 15 possibilites for what its primary key could be set to (e.g., ACE).

i. [T] We know that table **Q** has only one candidate key. If **Q** is in 3NF, it is also in BCNF.

j. [F] We know that table **Q** has only one candidate key. **Q** is in 2NF.

3. **Reverse Engineering.** *I've always wanted to direct.* (10 points) [Short Answer]

Person(sin, *name*, gender, studio, title, year)
FK (studio) references **Studio**,
FK (title, year) references **Movie**
Category(cat)
Movie(title, year, length, *director*, *studio*, *cat*)
FK (director) references **Person** (sin),
FK (studio) references **Studio**,
FK (cat) references **Category**
Cast(sin, title, year, role)
FK (sin) references **Person**,
FK (title, year) references **Movie**
Studio(studio, city)

Figure 2: Movies and Casts.

A relational schema for movies and casts is given in Figure 2. The primary key for each relation is indicated by the underlined attributes, as sin in **Person** (and are, hence, not nullable). Attributes in italics are *not* nullable, as *director* in **Movie**. The foreign keys with each are written as FK ...

The attributes title and year in table **Person** indicate that person's favorite movie. The attribute job of **Person** is allowed two values: 'actor' or 'director'. Only actors are allowed to act in (**Cast**) movies. Only directors are allowed to direct (director). A movie is produced by a studio. A person (actor or director) can be *contracted* by a studio (but by at most one studio), as indicated by studio in **Person**.

Imagine an E-R diagram for which the relational schema in Figure 2 is a correct translation.

a. (2 points) Is **Movie** a weak entity on **Category**? Briefly, why or why not?

No. The key to **Category**, {*cat*}, *is not a sub-key of* **Movie**.

b. (2 points) Is **Cast** likely a relationship or an entity in the E-R diagram? Briefly support your answer.

It is likely a relationship because its key is the union of its foreign keys.

c. (2 points) How many relationships does the entity **Person** participate in?

Four.

- *studio the person is with*
- *person's favourite movie*
- **director** *(in* **Movie***)*
- **sin** *(for actor cast) in* **Cast**

title	**year**	**length**	**director**	**studio**	**cat**
The Dell Crim Mysteries	2001	185 min	212839144	Columbia	mystery
The Dell Crim Mysteries	2000	82 min	212839144	Paramount	scifi

d. (2 points) Is it possible that the table **Movie** contain the rows as above?

Yes. The key for **Movie** *is* $\{title, year\}$*. The title value is the same for these two rows, but their year values are different.*

e. (2 points) Can we represent that the same actor (person) has several roles in a movie? Why or why not?

No. Cast can associate a given movie and person once, given the key of the table. That one row has a role value. There is no way to associate more than one role value.

4. **Normalization.** *Who's to say what's normal?* (10 pts) [Short Answer / Exercise]

1NF: Domain of each attribute is an *elementary* type; that is, not a *set* or a *record structure*.

2NF: Whenever $\mathcal{X} \mapsto A$ is a functional dependency that holds in relation **R** and $A \notin \mathcal{X}$, then either

- A is *prime*, or
- $\mathcal{X}$ is not a proper subset of any key for **R**.

3NF: Whenever $\mathcal{X} \mapsto A$ is a functional dependency that holds in relation **R** and $A \notin \mathcal{X}$, then either

- A is *prime*, or
- $\mathcal{X}$ is a key or a super-key for **R**.

BCNF: Whenever $\mathcal{X} \mapsto A$ is a functional dependency that holds in relation **R** and $A \notin \mathcal{X}$, then

- $\mathcal{X}$ is a key or a super-key for **R**.

An attribute A is called *prime* if A is in any of the candidate keys.

Figure 3: The Normal Forms.

a. (4 points) The relation ITYPL has the following functional dependencies.

$\mathsf{I} \mapsto \mathsf{TYPL}$. $\mathsf{T} \mapsto \mathsf{L}$. $\mathsf{TYP} \mapsto \mathsf{I}$.

Show the relations's candidate keys, and explain whether this is in BCNF, 3NF, and 2NF.

Clearly I is a candidate key. Thus TYP is also.
L is not prime and T ⊂ TYP, thus T ↦ L violates 2NF.
So this is not in 2NF, not in 3NF, and not in BCNF.

b. (4 points) The relation ITYPLA has the following functional dependencies.

$\mathsf{I} \mapsto \mathsf{TYPLA}$. $\mathsf{TYPL} \mapsto \mathsf{I}$. $\mathsf{TYA} \mapsto \mathsf{P}$.

Show the relation's candidate keys, and explain whether this is in BCNF, 3NF, and 2NF.

I is a candidate key. TYPL too. So is TYAL then, since TYA ↦ P.
Thus all the attributes are prime. This means the relation must be in 2NF and in 3NF.
However, it is not in BCNF. TYA ↦ P violates this since TYA is not a key nor superkey.

c. (2 points) Devise a *simple* example of a relation and functional dependencies such that it is in 2NF but not in 3NF.

$$A \mapsto B$$
$$B \mapsto C$$

So {A} is the only candidate key for ABC. The second functional dependency (FD) breaks 3NF since B is not a key or super-key and *C is not prime.*
This is in 2NF, though. The first FD is fine: {A} is key. The second FD is fine: {B} is not a sub-key.

Extra space.

Extra space.

Relax. Turn in your exam. Go home.