CSC 173
Homework 1
12th September, 2006
Due in class on Thursday, September 14th, 2006

Title (T)	Author (A)	Genre (G)	NumCopies (N)
Solaris	Lem	SF	2
Ender's Game	Card	SF	2
Ender's Shadow	Card	SF	1
Once and Future King	White	F	1
The Da Vinci Code	Brown	F	3
Best Karate	Nakayama	NF	1
Leonardo da Vinci	Arasse	NF	1
Leonardo da Vinci	Marani	NF	1
Name of the Rose	Eco	F	2

UserName (U)	Title (T)	Author (A)
ma	Ender's Game	Card
sd	Leonardo da Vinci	Arasse
sd	The Da Vinci Code	Brown
as	Name of the Rose	Eco
xz	Ender's Game	Card

UserName (U)	First (F)	Last (L)
ma	Max	Abernethy
sd	Sandhya	Dwarkadas
fm	Fabrizio	Morbini
as	Arrvindh	Shriraman
xz	Xiao	Zhang

The tables above represent a library database, consisting of a relation each for the books owned by the library, the books currently checked out by users, and information on each legitimate user of the library.

1. What attribute(s) constitute possible keys for each of the three given relations? Clearly state/explain any assumptions you must make.

2. What relation results from doing a natural join of the relation U-F-L with U-T-A[1]?

3. What relation is obtained by selecting from the relation T-A-G-N where (Author == Card OR Genre == SF)?

4. What relation results from doing a projection of the relation T-A-G-N onto the attributes Author and Genre?

5. If the T-A-G-N relation is stored in a hashtable where the Author's last name is the key of the hash table, what is the running time of a query for all book titles by a specific author? What is the running time of a query requesting all book titles of which there are two or more copies?

6. How would you determine how many copies of the book titled *The Da Vinci Code* by Brown remain in the library? I'm not interested in the answer — I'm interested in the query you would supply to the database to get the answer.

7. Trivia Assignment for Project 1 (to be turned in by sending e-mail to `cs173@cs.rochester.edu by noon on Friday September 15th)`:

   - What is the shortest, simplest Java program you can come up with that will repeatedly read lines of text from Unix standard input and write them to Unix standard output? (Send actual code. Test it first!)
   - Assuming that your `relation` class supports the `Enumeration` interface, give an example of a `for` loop that performs an operation on every tuple of a relation `R`.

[1] In all cases, I'm looking for the entire relation, i.e., all attributes (columns) and rows (entries) in the table, not just the scheme.