CSE 562 Final Exam *Solutions*

May 12, 2014

Question	Points Possible	Points Earned
A.1	7	
A.2	7	
A.3	6	
B.1	10	
B.2	10	
C.1	10	
C.2	10	
D.1	10	
D.2	10	
E	20	
Bonus	5	
Total	105	

Relational Algebra Operator Reference

Selection	$\sigma_c(R)$	c : The selection condition
Projection	$\pi_{e_1,e_2,\ldots}(R)$	e_i : The column or expression to project
Cartesian Product	$R_1 \times R_2$	
Join	$R_1 \bowtie_c R_2$	c : the join condition
Aggregate	$\pi_{gb_1,gb_2,\ldots,\texttt{SUM}(e_1),\ldots}(R)$	gb_i : group by columns, e_i : expression
Set Difference	$R_1 - R_2$	
Union	$R_1 \cup R_2$	

Starfleet Cargo Schema

```
CREATE TABLE Ship (
  ShipID int, Name string, Origin string, PRIMARY KEY (ShipID)
);
CREATE TABLE Manifest (
  ManifestID int, ShipID int FOREIGN KEY REFERENCES Ship(ShipID),
  LoggedDate decimal, SigningOfficer string,
  PRIMARY KEY (ManifestID), UNIQUE (ShipID, LoggedDate)
);
CREATE TABLE Cargo (
  ManifestID int FOREIGN KEY REFERENCES Manifest(ManifestID),
  RowID int, CargoDescription string, Value decimal,
  PRIMARY KEY (ManifestID, RowID)
)
```

The Tribble Query

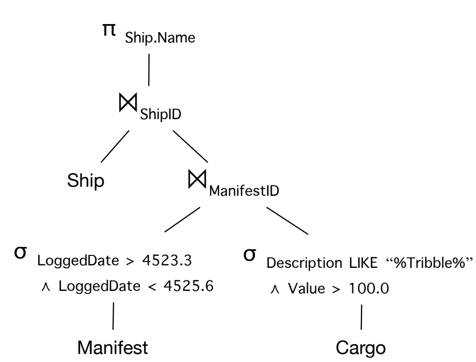

Part A. SQL Relational Algebra (20 points)

This question references the Starfleet Cargo Schema on Page 2.

1. (10 points) Write a SQL query that returns all ships with an item of cargo in its manifest that has a higher value than any item with an identical description.

2. (10 points) Write a SQL query that returns every port of origin (Ship.Origin) that does **not** have a ship that ever carried an item with the word "Tribble" in its description string.

3. As discussed in class, the right-outer join $⟖$ operates like a normal join, but producing an output for every tuple in the right-hand-side relation that is not joined by any tuples on the left. Every row created in this way has NULL values for left-hand-side attributes. Consider The Tribble Query on Page 2 with every join replaced by a right-outer join. Is the following query equivalent?

$$\pi_{Name}(\sigma_{Descr\ LIKE\ '\%Tribbles\%' \wedge Value > 100.0}(Ship ⟖ (\sigma_{LoggedDate}(Manifest ⟖ Cargo)))$$

If so, provide a proof using relational algebra equivalencies. If not, provide a counter-example.

Part B. Cost Estimation (20 points)

This question references the Tribble Query and schema on Page 2. Starfleet is considering plans for upgrading its data warehousing system, unifying ship registries and cargo manifests stored in databases across three different planets. Data is *partitioned* across three planets as shown in the chart below. Ship and Manifest records are partitioned independently. Cargo records are stored with their respective manifest.

Planet	# of Ship Records	# of Manifest Records	# of Cargo Records
Earth	100	10,000	50,000 (5 cargo / manifest)
Vulcan	500	6,000	12,000 (2 cargo / manifest)
Andoria	0	20,000	200,000 (10 cargo / manifest)

- Earth has a Clustered B+Tree index on $Manifest.Date$ and $Cargo.Value$, and a Clustered Hash index on $Ship.ID$.
- Vulcan has an Clustered B+Tree index on $Cargo.Value$ and a Unclustered B+Tree index on $Manifest.Date$, a Clustered Hash index on $Manifest.ID$ and $Ship.ID$.
- Andoria has a Inverted Index (which allows lookups on LIKE) on $Manifest.Description$, and Clustered Hash indexes on $Manifest.ID$ and $Ship.ID$.

Assume that one page holds 100 tuples (of any relation). Assume that a bloom filter consumes 1 page per 1,000 keys encoded, and that bloom filters have a 20% overhead from false positives. Each B+Tree index has 3 levels, and Hash indexes are sufficiently large to have no overflow pages. Assume that each site has sufficient memory to keep the entire query in memory. Assume that each ship has an average of 60 manifests, that the selection predicate on *LoggedDate* eliminates 90% of tuples, and that the selection predicates on *Description* and $Value$ eliminate 99% and 50% of tuples, respectively.

For each of the following parallel plans, state (1) The optimal access paths at each site (optimal by number of Page IOs), (2) The number of page IOs required at each site, and (3) The number of pages shipped between sites.

1. (10 points) Compute the join of Manifest and Cargo locally at each site. Send the results, and Vulcan and Andoria's partions of the Ship relation to Earth, and compute the final result there.

2. (10 points) Compute the join of Manifest and Cargo locally at each site. Have each site use bloom join to compute the join with Ship. The final join result is shipped to Earth.

Part C. Transactions (20 points)

A new database application: TrekDB is considering using a new form of serializability: Warp Serializability. Warp Serializability acts like View Serializability, but also maintains sufficient metadata with each write to allow certain kinds of write operations to be reordered. Write operations belonging to a given "Warp Group" can be safely reordered, even if applied to the same object. The syntax `W(A,1)` indicates a write to object `A`, in warp group `1`.

For each of the following schedules, and each of Conflict-, View-, and Warp-Serializability:

- State whether the schedule is X-serializable (i.e. serializable under the given equivalence).
- If the schedule is X-serializable, provide an X-equivalent serial schedule.
- If the schedule is not X-serializable, identify a minimal subset of the schedule that prevents it from being X-serializable.

1. `Xact 1:R(A), Xact 2:W(A,1), Xact 2:W(B,2), Xact 3:W(C,3), Xact 1:W(C,3), Xact 3:W(B,3)`
   (a) **Conflict**:

   (b) **View**:

   (c) **Warp**:

2. `Xact 4:W(D,3), Xact 1:R(A), Xact 2:W(A,1), Xact 2:W(B,2), Xact 3:W(C,3), Xact 1:W(C,3), Xact 3:W(B,3), Xact 4:W(C,4)`
   (a) **Conflict**:

   (b) **View**:

   (c) **Warp**:

Part D. Logging and Recovery (20 points)

A database crashes with the following state written to disk.

Page	Data	Last Write
1	A = 1 B = 10 C = 26	TS 5
2	D = 51 E = 47 F = 38	TS 1
3	G = 50 H = 20 J = 65	TS 7

and the following log

Timestamp	1	2	3	4	5
Transaction	Xact 3	Xact 1	Xact 3		Xact 5
Object/Op	E	H	COMMIT	CHECKPOINT	A
New Value	47	19			1
Old Value	98	88			17

Timestamp	6	7	8	9	10
Transaction	Xact 6	Xact 5	Xact 6	Xact 1	Xact 7
Object/Op	F	H	COMMIT	COMMIT	A
New Value	99	20			55
Old Value	39	19			1

The checkpoint at timestamp 4 stores two transactions: Xact1 with a prior write at TS 2, and Xact2 with no prior writes.

Replay the three phases of the ARIES recovery protocol step by step. Be sure to include every modification applied by the protocol to the transaction table, data pages, and log.

Part E. True/False (25 points)

True or False. Mark one answer. Any explanatory text will be ignored.

1. A declarative language like SQL is designed to force you to specify *how* you want to compute what you want.
   True | **False**
2. $R \bowtie S$ is equivalent to $R \times S$ if the schemas of R and S have two attributes in common.
   True | **False**
3. Two selection predicates have preserve (have selectivities) 10% and 20% of the data respectively. This guarantees that 2% of records will satisfy both predicates.
   True | **False**
4. Intersection can not be synthesized from the 4 core relational algebra operators: Select, Project, Cross Product, Union.
   True | **False**
5. Strict 2-Phase locking is guaranteed to produce a conflict serializable schedule.
   True | **False**
6. Any conflict-serializable schedule is can be produced by strict 2-phase locking.
   True | **False**
7. All view serializable schedules are conflict serializable.
   True | **False**
8. All serializable schedules are view serializable.
   True | **False**
9. In a shared-nothing system, multiple CPUs are attached to an interconnect. Each CPU has its own memory, and all access the same shared array of disks.
   True | **False**
10. Consider the relation $R(A, B, C)$. If there is a functional dependency from A to (B, C), then (A, B) is a candidate key.
    True | **False**
11. Write Ahead Logging is sufficient to recover from a crash.
    True | **False**
12. A star schema involves a fact table and one or more dimension tables.
    True | **False**

13. The cube operator can be synthesized from n^2 regular aggregate queries, where n is the number of group-by attributes in the cube.
True | **False**

14. Column stores lay out entire tuples sequentially on disk/in memory.
True | **False**

15. In Map/Reduce, the Map task can be carried out with per-tuple parallelism, and the reduce task can be carried out with per-group parallelism.
True | **False**

16. Nested Relational Algebra is standard Relational Algebra plus a nest and an unnest operator.
True | **False**

17. The following monad algebra expression computes the intersection of two sets
$flatten \circ map[ITE] \circ flatten \circ map[pairwith \circ \langle \pi_{right}, \pi_{left} \rangle] \circ pairwith$
where ITE is $if\ (\pi_{left}\ id) = (\pi_{right}\ id)\ then\ \{(\pi_{left}\ id)\}\ else\ \emptyset$
True | **False**

18. Every statement in XPath can be expressed in XQuery
True | **False**

19. Every statement in XQuery can be expressed in XPath
True | **False**

20. Any given NoSQL database can handle all workloads better than a One-Size-Fits-All Database
True | **False**

Bonus Question (5 pts): The TPC-H Schema uses the singular for nearly all of its table names: Customer, Lineitem, Part. Only Orders is pluralized. Why?