

**CSCI 3357: Database System Implementation**  
Homework Assignment 10  
Due Thursday, November 30

In HW 8 you modified the query processor to implement the *union* and *rename* relational algebra operators. In this assignment you will extend these modifications into the parser and planner. Warning: This is a time-consuming assignment; start early.

If you do not have correct versions of `RenameScan` and `UnionScan`, use mine from the posted HW 8 solutions.

1. Write classes `UnionPlan` and `RenamePlan`. I care mostly about their `open` and `schema` methods. Make reasonable assumptions about how to implement their three statistical methods. It really doesn't matter what values you use, but I want to see evidence that you know what these methods are for.

2. Standard SQL uses the `UNION` keyword to connect select statements (much in the same way that the `AND` keyword separates terms in a predicate). For example, the following query returns the names of students having `sid` 1, 3, or 5:

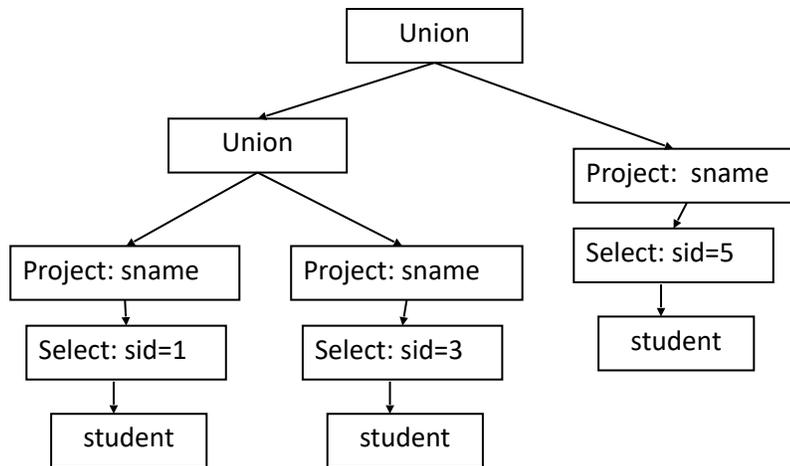
```
select sname from student where sid = 1
union
select sname from student where sid = 3
union
select sname from student where sid = 5
```

You can incorporate this new functionality into the SimpleDB grammar by adding the following rule:

```
<UnionQuery> := <Query> [ UNION <UnionQuery> ]
```

a) Modify the classes `Lexer` and `Parser` to implement this new rule. The parser method for the syntactic category `<UnionQuery>` should return a list of `QueryData` objects—one for each subquery.

b) Modify `BasicQueryPlanner` to have a method `createUnionPlan`, which takes a list of `QueryData` objects as its argument and creates a query tree having a `UnionPlan` node for each `UNION` keyword in the query. If there are no `UNION` keywords in the query then the list will contain one `QueryData` object and its plan will be the same as the plan that would be created by the current basic query planner. For example, the plan for the above query should correspond to the following query tree:



c) Rename the `createPlan` method of `QueryPlanner` so that its name is `createUnionPlan` and modify its first argument to be `List<QueryData>`. Also modify the `createQueryPlan` method of `Planner` so that it calls `parser.unionQuery` and `qplanner.createUnionPlan` instead of `parser.query` and `qplanner.createPlan`. You will of course also need to adjust its variables of type `QueryData` to be `List<QueryData>`.

3. Standard SQL uses the `AS` keyword in its select clause to rename a field. For example, the following query returns the id, name and major of the students graduating in 2020, with the field `Sid` renamed as `StudentNum` and the field `MajorId` renamed as `MajorDept`:

```

select Sid as StudentNum, SName, MajorId as MajorDept
from Student
where GradYear = 2020

```

You can incorporate this functionality into the SimpleDB grammar by modifying the rule for `<SelectList>` and adding a rule for the new category `<SelectField>`, as follows:

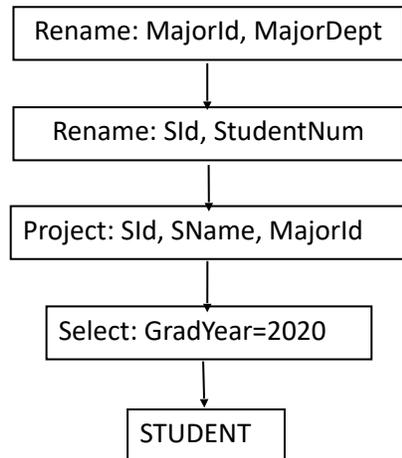
```

<SelectList> := <SelectField> [ , <SelectList> ]
<SelectField> := <Field> [ AS <Field> ]

```

a) Modify the `Parser` class to implement these grammar changes. Create a class `FieldData` to hold the values extracted from the `<SelectField>` rule. Then modify `QueryData` so that its variable `fields` is a list of `FieldData` objects instead of a list of strings.

b) Modify step 4 of the `BasicQueryPlanner` method `createPlan` to add `RenamePlan` nodes to the query plan it creates. There should be one node for each renamed field, and these nodes should appear (in any order) above the `ProjectPlan` node. For example, the plan for the above query should correspond to the following query tree:



Note that the renaming happens after the project operation occurs, which means that the where-clause can only refer to original field names, before the renaming occurs. For example, consider the query below. Note that the predicate contains the term `SId > 5` and not `StudentNum > 5`.

```
select SId as StudentNum, SName, MajorId as MajorDept
from Student
where GradYear = 2020 and SId > 5
```

One consequence of changing the planner is that all existing planner files will need to be updated (such as `HeuristicQueryPlanner.java` and `BetterQueryPlanner.java`). You don't need to do this. When I was debugging my code, I chose to delete these files from the project because I knew they weren't used. Feel free to do the same.

---

Once you get everything to work, have some fun. Run the `SimpleJ` client program, and execute some SQL commands that involve union and renaming. For example, try this:

```
select sname as person from student
union select prof as person from section
```

To help with debugging, you can download my file `HW10Test.java`. I don't guarantee that it detects all bugs.

When you are done, create a zip file containing the nine files you created or modified, namely `UnionPlan`, `RenamePlan`, `Parser`, `Lexer`, `QueryData`, `FieldData`, `BasicQueryPlanner`, `QueryPlanner`, and `Planner`. Then submit this zip file to Canvas.