

Exercise 1

As we already know, the number of non-symmetric csg-cmp-pairs ($\#ccp$) depends on the query graph:

$$\#ccp^{chain}(n) = 1/6 * (n^3 - n)$$

$$\#ccp^{cycle}(n) = (n^3 - 2n^2 + n)/2$$

$$\#ccp^{star}(n) = (n - 1)2^{n-2}$$

$$\#ccp^{clique}(n) = (3^n - 2^{n+1} + 1)/2$$

Luckily, we don't have to store all of them in our DP-table. For all DP-based algorithm that don't consider cross products, we store only the cheapest plan (seen so far) for each connected sub graph.

$$\#csg^{chain}(n) = n(n + 1)/2$$

$$\#csg^{cycle}(n) = n^2 - n + 1$$

$$\#csg^{star}(n) = 2^{n-1} + n - 1$$

$$\#csg^{clique}(n) = 2^n - 1$$

Exercise 1 a)

For a star query with with $n = 20$ relations, how many plans do you have to store in your DP-table? What about $n = 30$?

Exercise 1 b)

If each plan consumes 40 bytes of memory. Then how much memory consumes the DP-table for 30 relations?

Exercise 1 c)

How to approach large problem sizes?

Exercise 2

Exercise 2 a)

Recall the introductory DP exercise:

Walking up the stairs. How many steps can you take at a time? Let's say up to three! Then how many ways are there to walk up a staircase with n steps?

... This time, use memoization to answer the question!

Exercise 2 b)

Implement `MemoizationJoinOrdering`. You may use the helper classes provided in the solution code.

Exercise 3

Modify `MemoizationJoinOrdering` such that cross products are excluded.

Exercise 3 a)

What do you observe with regard to the connection tests? Compare this to `DPsub`.

Exercise 3 b)

Name Pros and Cons compared to DP/ bottom-up approaches.