

Exercise 5

Released: May 06, 2019 · Discussion: May 13, 2019

1 JOIN Implementation

JOIN operators are one of the most important operators in DBMS. In this exercise you have to implement at least two different **JOIN** algorithms in a given template. Add an implementation for:

- a *Nested Loops Join*
- an *Index Nested Loops Join*

Download the file `03_join.zip` from the course website¹ and extract it. If you extract it into the folder of the previous projects you may have to merge the `CMakeLists.txt` files. You have to complete the files `src/operators/nested_loops_join.cpp` and `src/operators/index_nested_loops_join.cpp`.

1. Which implementation is faster?
2. What happens if the inner and the outer relation are switched? Describe the observation and explain it.

Build instructions:

1. Extract the archive and navigate into the extracted folder.
2. Run `cmake` to create a makefile for your system: `cmake .`
3. Run `make` to create an executable binary file: `make`
4. Execute the created binary file (e.g. `./03_join` on linux)

¹http://dbis.cs.tu-dortmund.de/cms/en/teaching/ss19/arch-dbms/exercises/03_join.zip