

Worksheet 8

Part 1: Physical Plan Analysis

Compute the cost of each physical plan below. For each select, choose the operation that would minimize the cost the most. You have the following data:

Places(<u>location</u>) Symptoms(<u>symptom</u>) SampleDNA(<u>sid</u> , symptom, sequence, located) -- located is a foreign key to Places -- symptom is a foreign key to Symptoms ExposedPlace(toxin, <u>symptom</u> , located) -- located is a foreign key to Places -- symptom is a foreign key to Symptoms	M = 100 Unclustered index (located) for SampleDNA Unclustered index (toxin) for ExposedPlace Unclustered index (symptom) for ExposedPlace	
B(SampleDNA) = 10^5 T(SampleDNA) = 10^6	B(ExposedPlace) = 10^4 T(ExposedPlace) = 10^5	B(Places) = 10^3 T(Places) = 10^6 B(Symptoms) = 10^1 T(Symptoms) = 10^4

Part 2: Distributed Computing Problem Solving

Assume for these problems that you have a relation SampleDNA(sid, symptomatic, sequence, located) Sequences are of, at most, length 2000. All sequences only contain the nucleotides A, T, G, and C.

SampleDNA			
sid	symptomatic	sequence	located
49396937	T	ATTTTCGATGCGCGTAAA...	Seattle WA
68478053	F	ATTCCGATGCGCGAAAA...	Boston MA
...	

For each of the problems write MapReduce pseudo-code and a Spark function to compute the result.

For MapReduce, assume you are given sid as a key and the remaining attributes as a tuple value.

For Spark, assume you are given an RDD r. Compute the information (don't worry about outputting).

1. Count number of symptomatic samples that were found in each location. Get the max count (you don't need the associated location).

2. For symptomatic samples find the distribution of the nucleotides (i.e. count the number of times nucleotide X appears in index n), for each nucleotide. Using that information determine which indexes have an occurrence of nucleotide T greater than 95%. (HW is not nearly this complicated)