

MapReduce

Use pseudocode to write MapReduce functions necessary to solve the problems below. Also, make sure to fill out the correct data types. Some tips:

- The input to each MapReduce job is given by the signature of the `map()` function.
- The function `emit(key k, value v)` outputs the key-value pair `(k, v)`
- You may use the `for(var in list)` syntax to iterate through `Iterables`, or you can call the `hasNext()` and `next()` functions
- Data types you may use are: `int`, `float`, `String`, list of these primitives, and custom data types composed of these primitives

1. Given a set of classes that students have taken, output each student's name & total GPA.

Declare any custom data types here: CourseData: int courseID float studentGrade //a number from 0-4	
map(String student, CourseData value):	reduce(_____key, Iterable< _____> values):

2. Compute the list of mutual friends between each pair of friends in a social network. Each person on the network is identified by a unique `int` ID. The `intersection(list1, list2)` method returns a list that is the intersection of `list1` and `list2`.

Declare any custom data types here:	
map(int personID, list<int> friendIDs):	reduce(_____key, Iterable< _____> values):

3. A. Given a set of coins and each coin's owner, compute the number of coins of each denomination that each person has.

Declare any custom data types here:	
map(String person, String coinType):	reduce(_____ key, Iterable< _____ > values):

- B. Using the output of the first MapReduce, compute the amount of money each person has. The function `valueOfCoin(String coinType)` returns a float corresponding to the dollar value of the coin.

map(_____ key, _____ value):	reduce(_____ key, Iterable< _____ > values):
-------------------------------------	---

Warehouse-Scale Computing

Power Usage Effectiveness (PUE) = (Total Building Power) / (IT Equipment Power)

Total Building Power = IT Equipment + Power supplies + Networking equipment + Cooling equipment

Sources speculate Google has over 1 million servers. Assume each of the 1 million servers draw an average of 200W, and that Google pays an average of 6 cents per kilowatt-hour for datacenter electricity.

- a) Estimate Google's annual power bill for its datacenters. Ignore the power cost of networking equipment. Assume 365 days (8760 hours) in a year.

- b) Google reduced the PUE of a 50,000 machine datacenter from 1.5 to 1.25 without decreasing the power supplied to the servers. What's the cost savings per year?