

Quiz 3:

A system of traffic cameras all over town continuously snaps stills. A system based on deep learning extracts automatically license plate information from vehicles in a still. The data is then stored in a very large distributed database. A record in this database will consist of location information, the license plate of the vehicle, the description of the vehicle (such as the make and color of a four-wheeler), the direction travelled, and a time stamp. You can imagine the size of the database by realizing that each trip with a vehicle will lead to hundreds records.

Assume that some drivers try to avoid automatic fines by making parts of the license plate unreadable. Fortunately for the municipal corporation, the license plate reading software is capable of reading partial license plates. To identify the scofflaw, the authorities are interested in finding routes a particular vehicle with a partially obscured license plate is traveling routinely. At a later step, this information is combined with a list of all vehicles matching the partial license and their description to identify the owner of the vehicle and then heavily fine them.

Your task is however to reconstruct routes giving a partial license plate and a description of the vehicle as extracted from traffic camera footage.

Set up a single map-reduce job that will allow you to reconstruct the route of a vehicle through the town.