

1 Distributed Databases: 2PC

An analogy for two-phase commit is getting married. Explain in detail how “getting married” maps to 2PC. (Note: this example was invented by Jim Gray who invented transactions.).

2 Yahoo!’s PNUTS

Yahoo!’s PNUTS adapts some known technology from standard database management systems. Identify at least **four** technologies that were presented before in this lecture. Give the corresponding terms in the PNUTS system, and explain how they are applied and modified there.

3 map/reduce

Assume you want to share the shopping load in your flat-sharing community. You are 5 people and you have a shopping list of 400 items from 14 different shops that need to be organized.

1. How would you model this task in the map/reduce framework? How do you implement map() and reduce()?
2. What if a majority of the items (> 50%) is available in a single shop? Would that change your implementation of map() and reduce()?
3. As an extension to the previous tasks: How would you make sure that everybody gets the same amount of shopping workload? (In the sense that everybody spends the same amount of time. Note: all 5 people doing all shopping together is not a valid answer...)