

Bayesian Personalized Ranking (BPR) with Spark

Alfredo Láinez Rodrigo *and* Luke de Oliveira



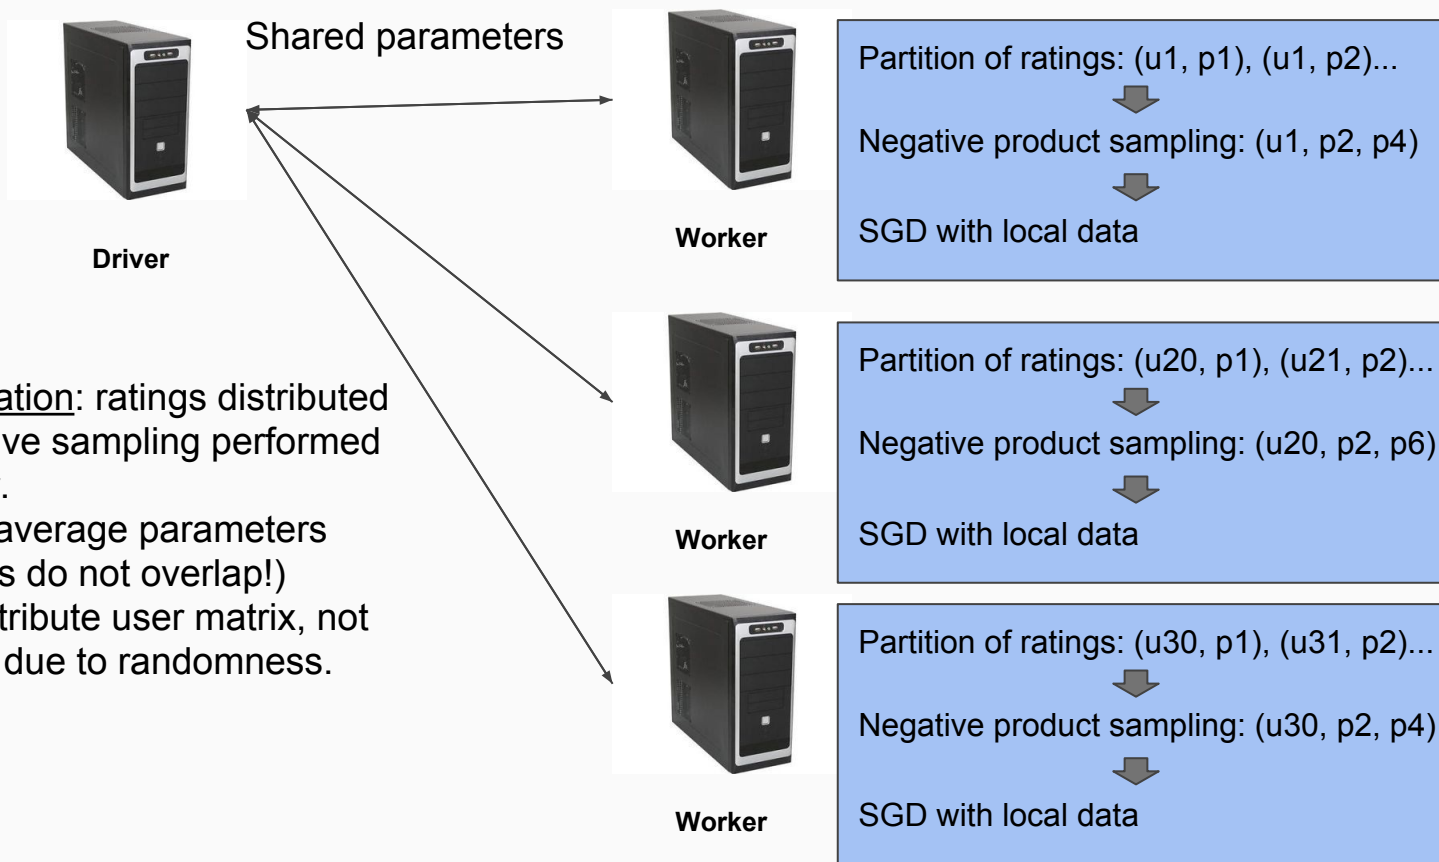
Bayesian Personalized Ranking for Implicit Feedback (BPR)

- Bayesian framework for personalized recommendations.
- Abstract away model from recommender system estimation

$$p(\Theta | \succ_u) \sim p(\succ_u | \Theta) p(\Theta)$$

- Implicit feedback \rightarrow maximize AUC w.r.t. negative samples
- BPR expressed as optimization problem \rightarrow SGD
- Learn ordering \succ_u over (i, j) items
- Chosen underlying model: Matrix Factorization

Parallelization in Spark



- Data parallelization: ratings distributed by user, negative sampling performed in each worker.
- Parallel SGD: average parameters (note that users do not overlap!)
- Possible to distribute user matrix, not product matrix due to randomness.