# A Lock-free Approach to Parallelizing Personalized PageRank Computations on GPU

**Zhigang WANG, Ning WANG, Jie NIE, Zhiqiang WEI, Yu GU, Ge YU**

Frontiers of Computer Science, DOI: 10.1007/s11704-022-1546-2

# Problems & Ideas

- Problems of parallelizing iterative PPR computations on GPU.
  - Existing “one-pass” design generates high concurrent criterion detections of vertex update, leading to heavy locking costs.
  - Existing push-based solution cannot avoid message write races issued by massive but different GPU threads.

- Ideas: a hybrid framework on top of newly designed lightweight push and lock-free pull models

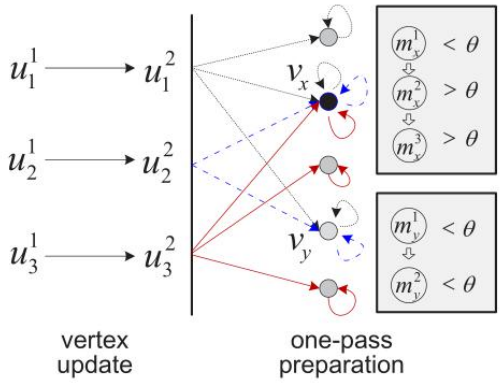


(a) Existing one-pass dataflow

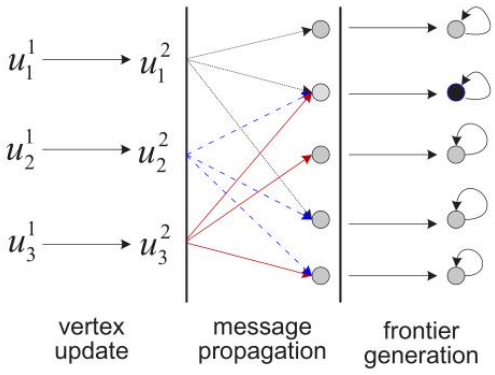


(b) Separated dataflow

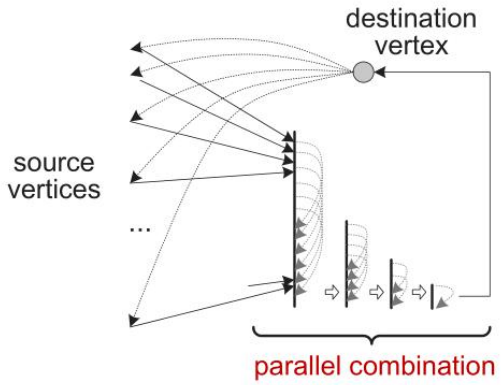


(c) Parallel binary combination

# Main Contributions

- Contributions:
  - Lightweight Forward Push (*pushopt*) by separating message propagation from criterion detection with a newly inserted global barrier
  - Lock-free Forward Pull (*pull*) with parallel combination optimization
  - Hybrid framework (*hybrid*) with ability of dynamically selecting a better model for a specific iteration

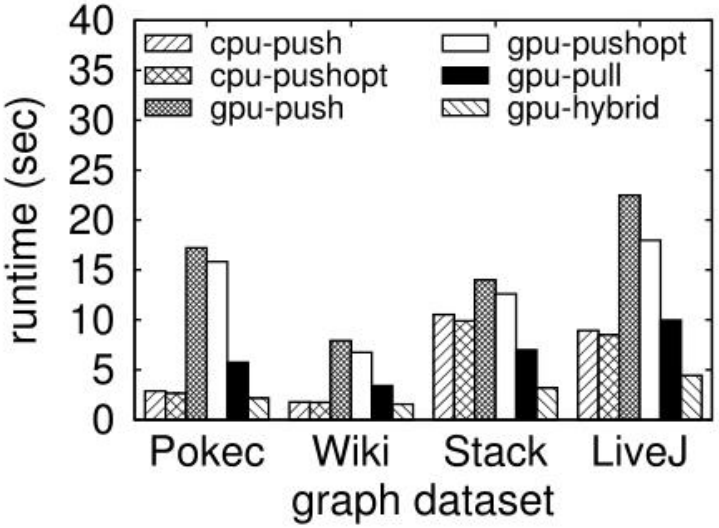


(a) Overall runtime comparison on CPU and GPU devices

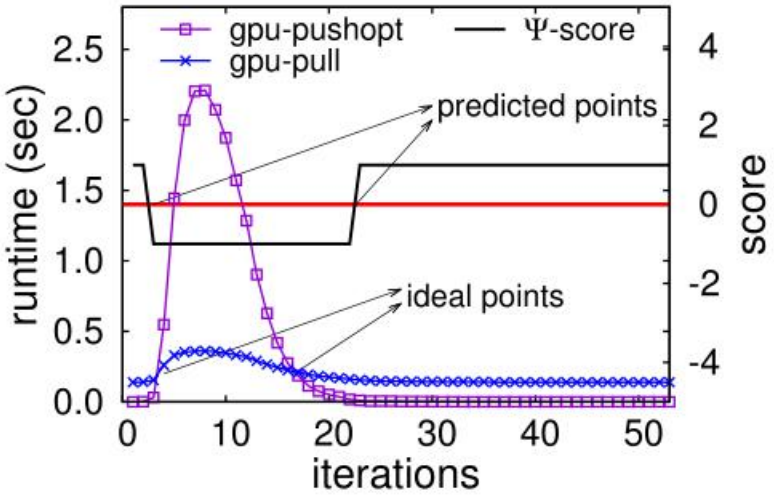


(b) How likely *hybrid* can seed the optimal model