Communication-Efficient Distributed SVD via Local Power Iterations

Xiang Li, Shusen Wang, Kun Chen, Zhihua Zhang
School of Mathematical Sciences, Peking University, China
Department of Computer Science, Stevens Institute of Technology, USA

Distributed Power Iteration (DPI)

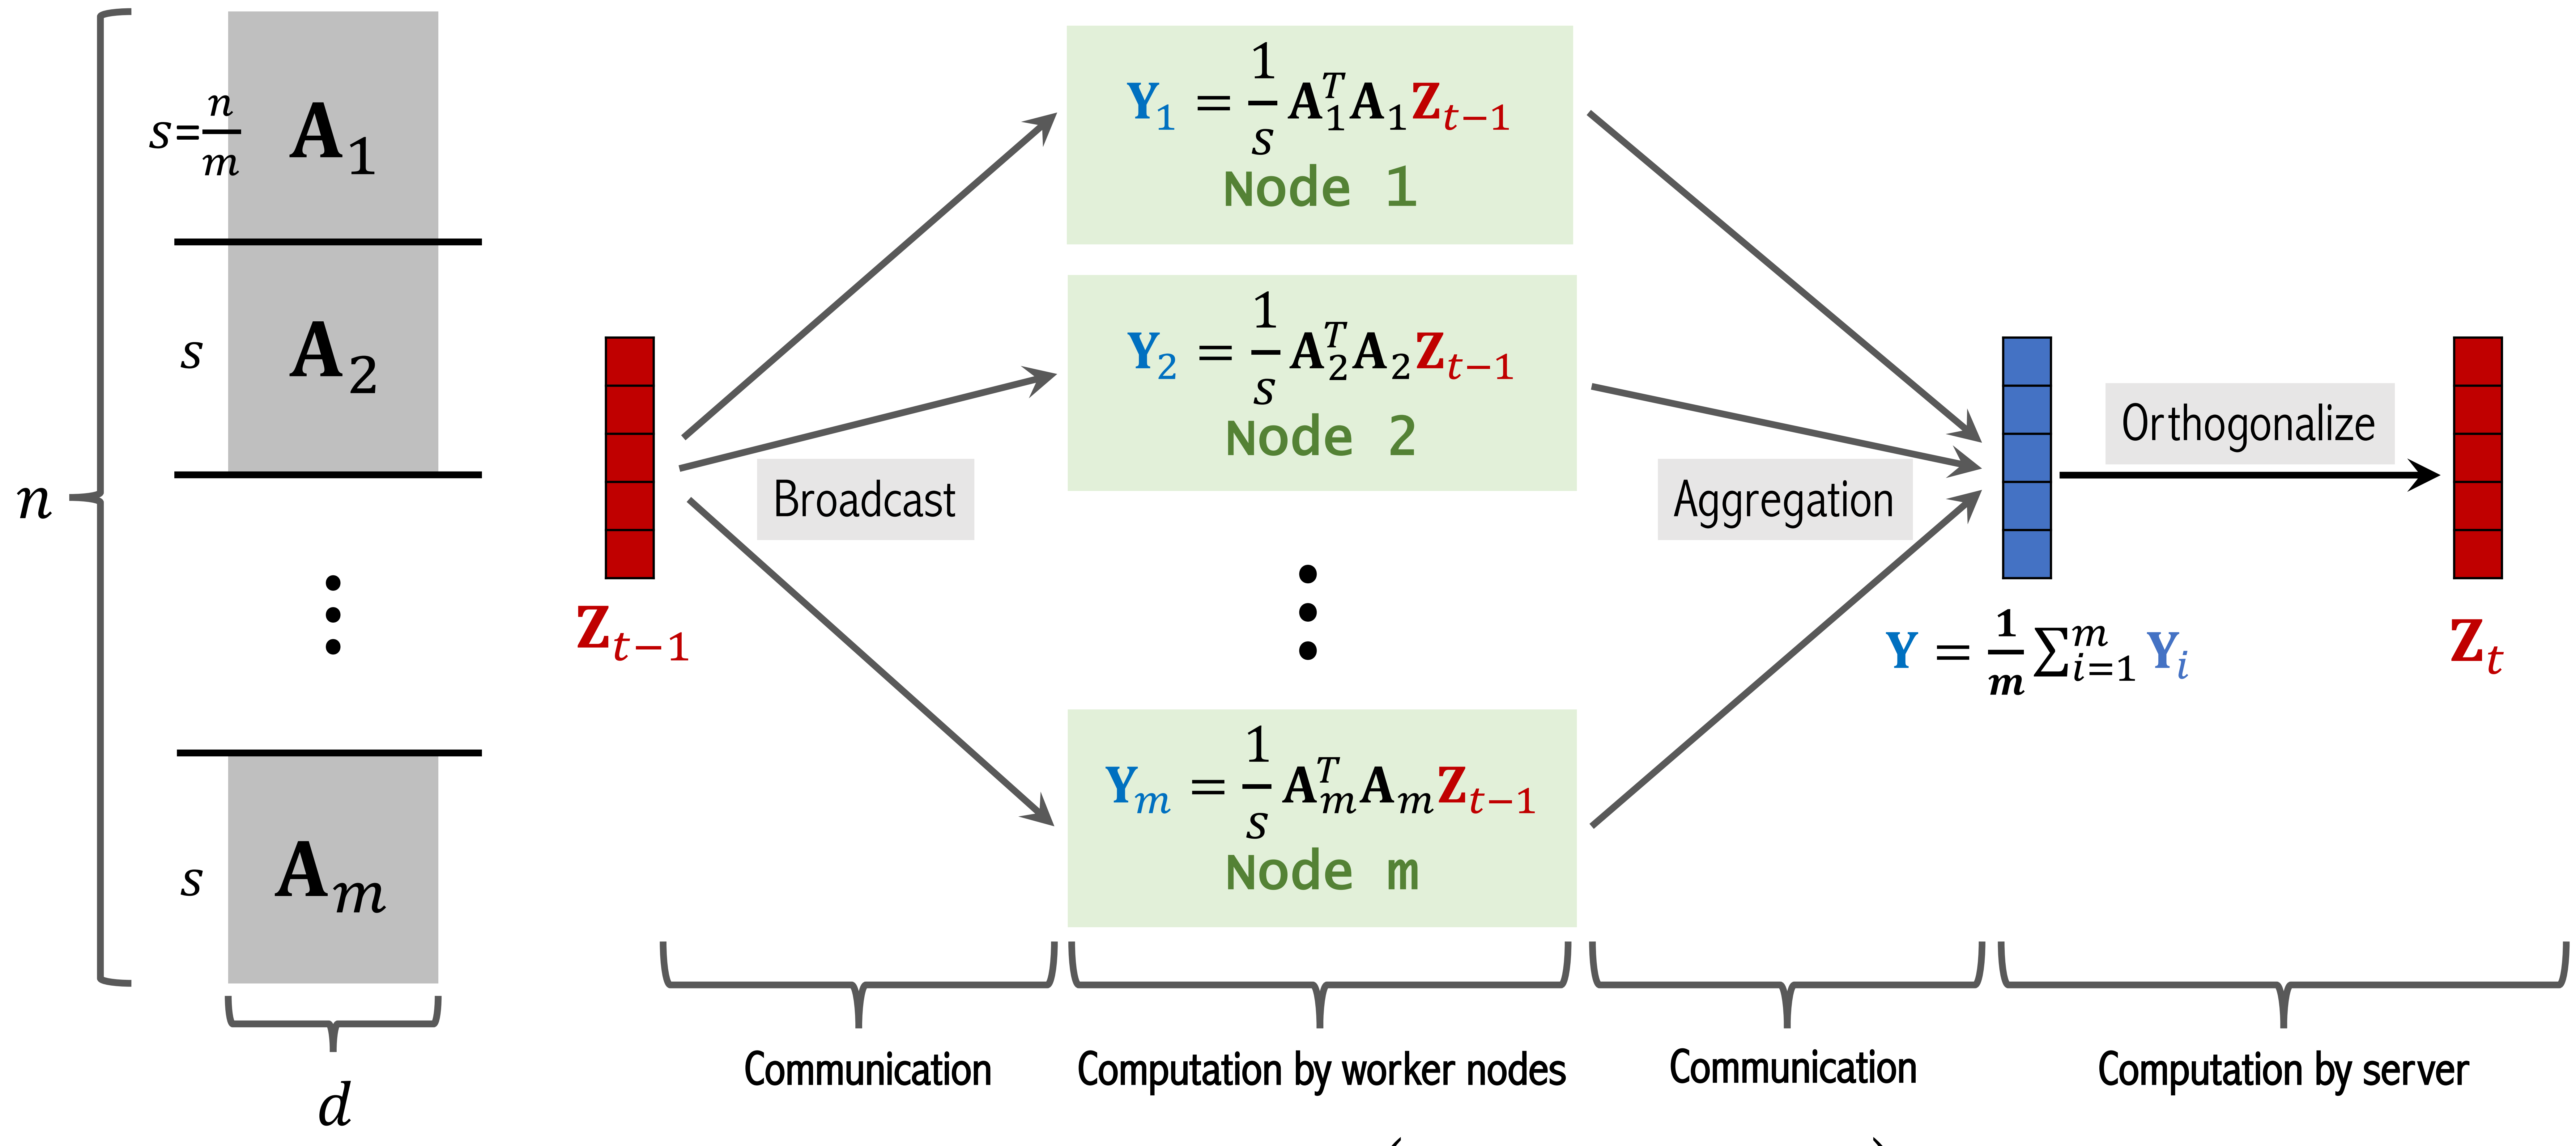

DPI communicates at each round and thus need $\Omega\left(\frac{\sigma_k}{\sigma_k - \sigma_{k+1}}\log\left(\frac{d}{\epsilon}\right)\right)$ **communication.**

Local Power (LP)

p local iterations

aggression weighted by $\{O_i\}$

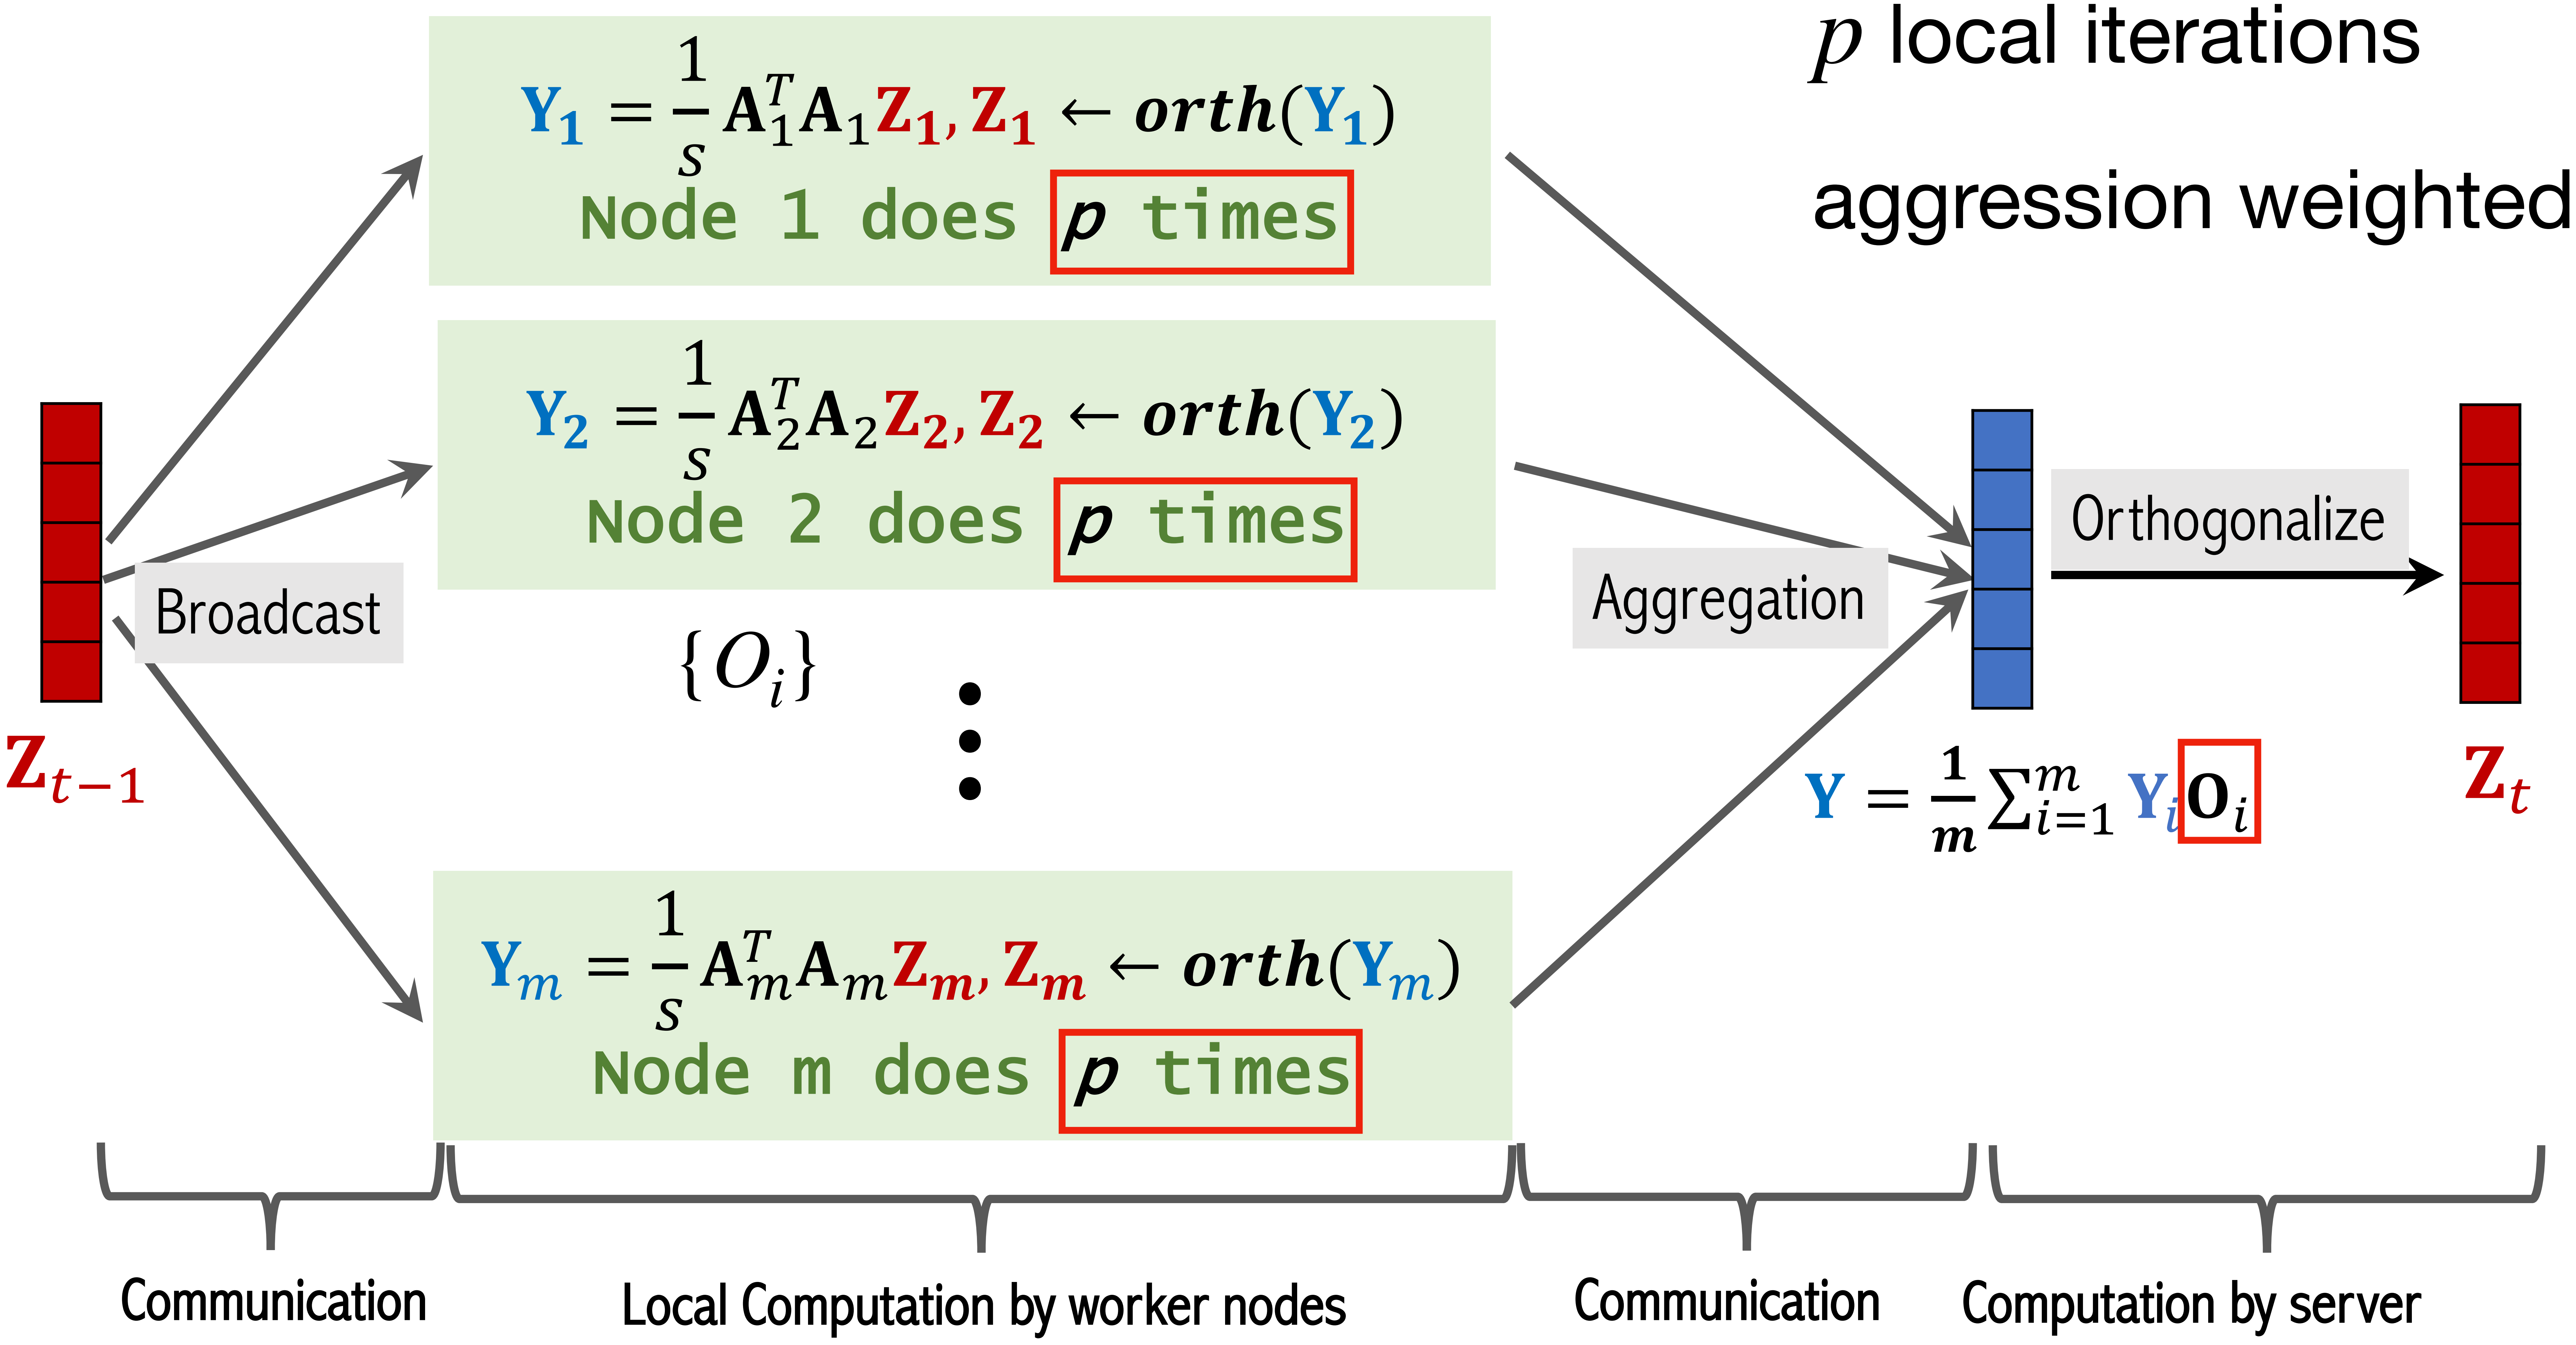

How to compute $\{O_i\}$?

Pure aggregation: $O_i = I_d$

Orthogonal procrustes transformation (OPT):

$$O_i = \arg \min_{\text{orth } O} \|Z_1 O - Z_i\|_F^2$$

Sign-fixing: use diagonal matrices with ± 1 diagonal entries:

$$O_i = \arg \min_{\text{diag } O \text{ with } \pm 1} \|Z_1 O - Z_i\|_F^2$$

The computation cost of sign-fixing is much smaller.

LP is Communication Efficient

- If local data matrices are similar enough or p is not too large, LocalPower save communication by a factor of p.
- When OPT is used, the restriction on local data matrices is much smaller than that of pure aggregation.
- OPT and sign-fixing are more stable than pure aggregation.
- Decaying p helps us better trade-off the communication efficiency and final error.

Thank You !