


Efficient Protocol for Collaborative Dictionary Learning in Decentralized Networks



T. Idé (井手 剛)¹, Raymond Rudy², and Dzung T. Phan¹

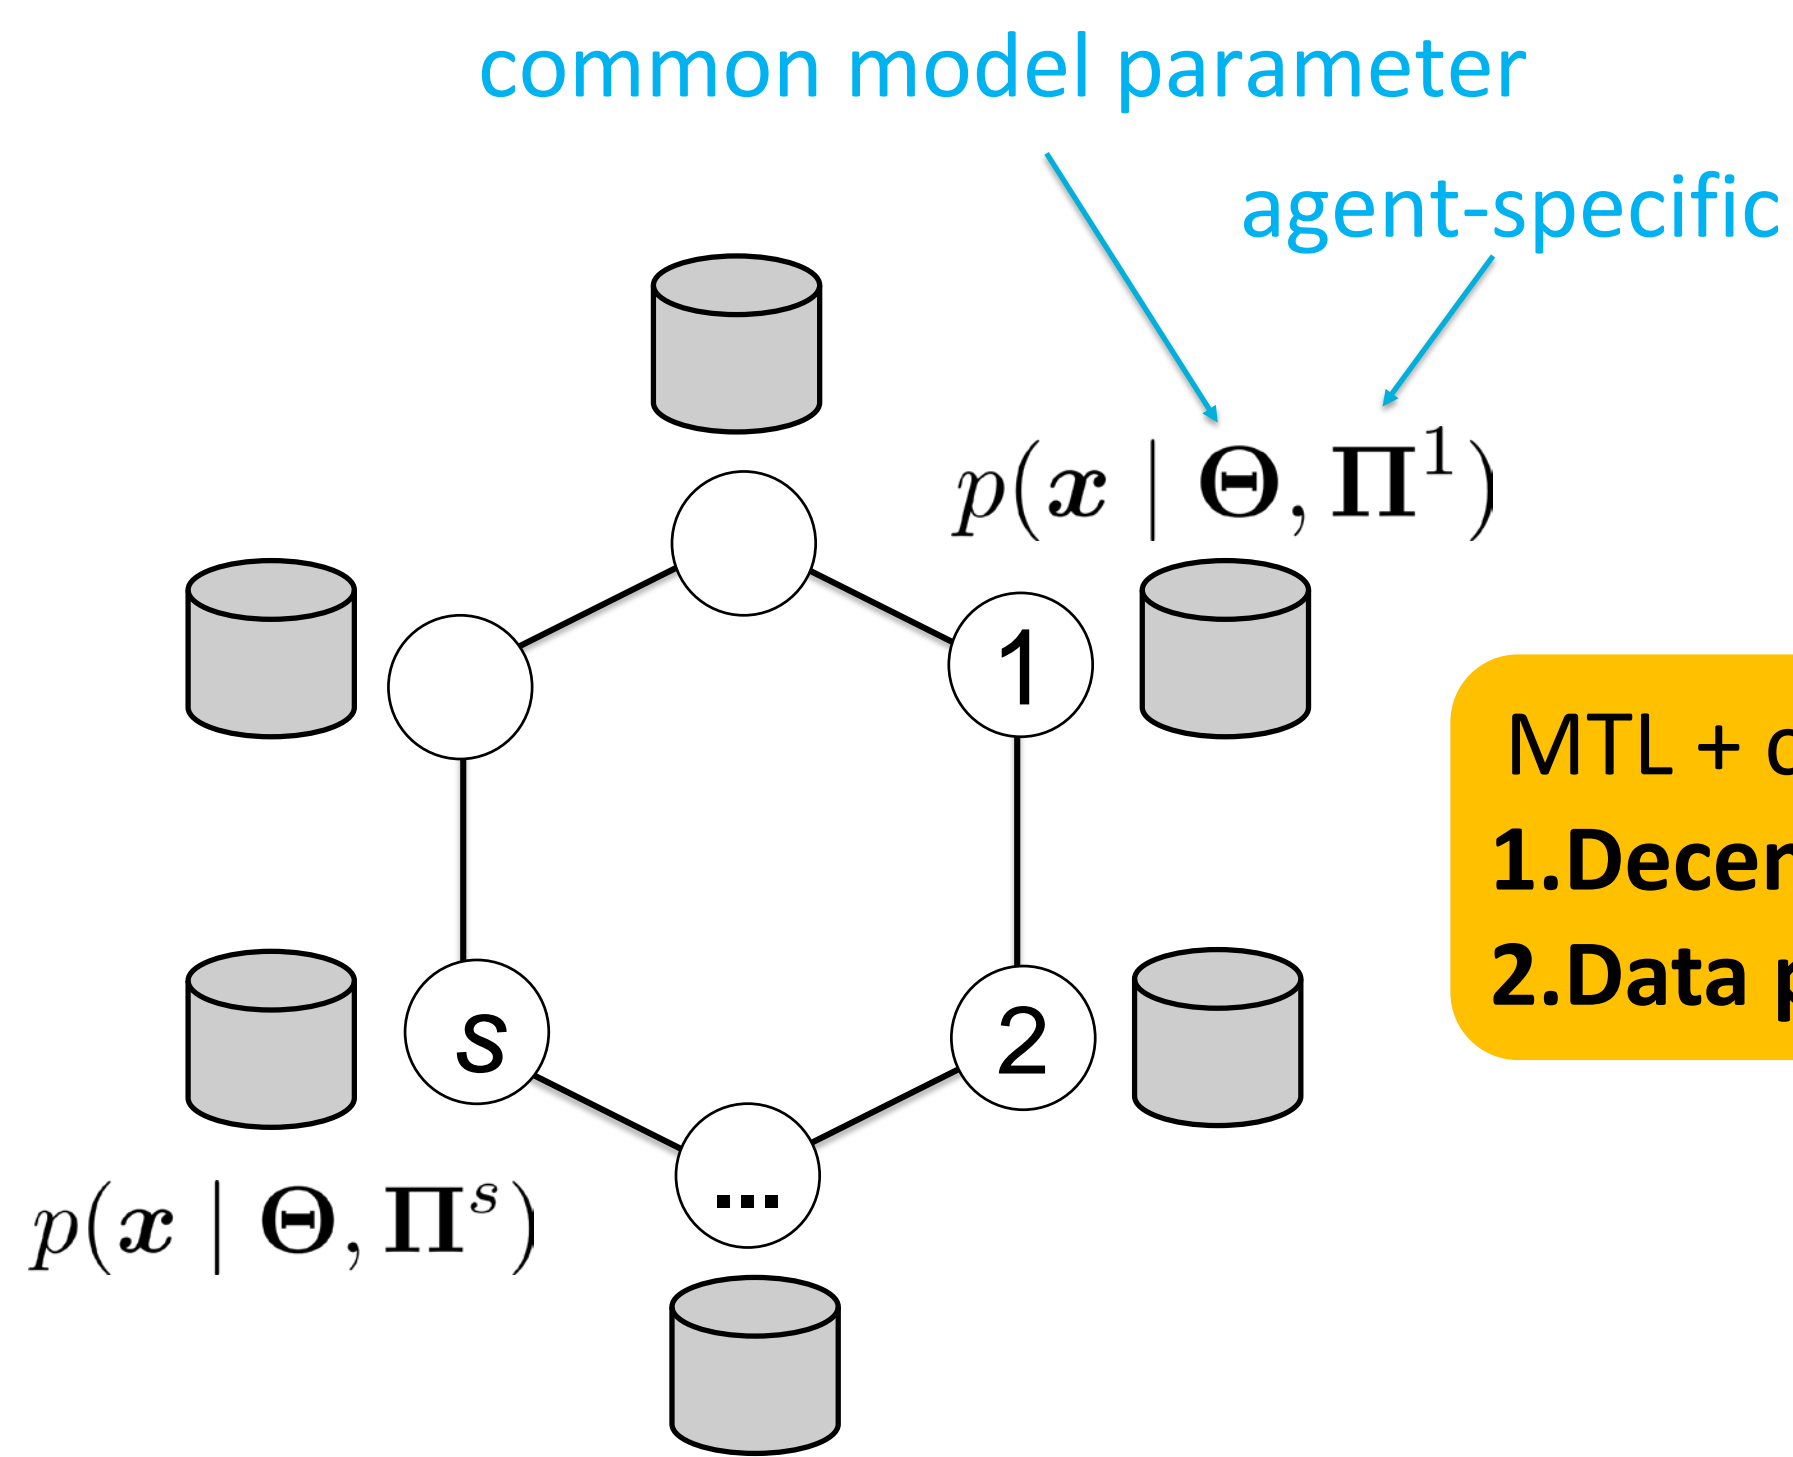
¹ IBM Research, T.J. Watson Research Center

² IBM Research - Tokyo



Problem setting

Decentralized multi-task density estimation with data privacy

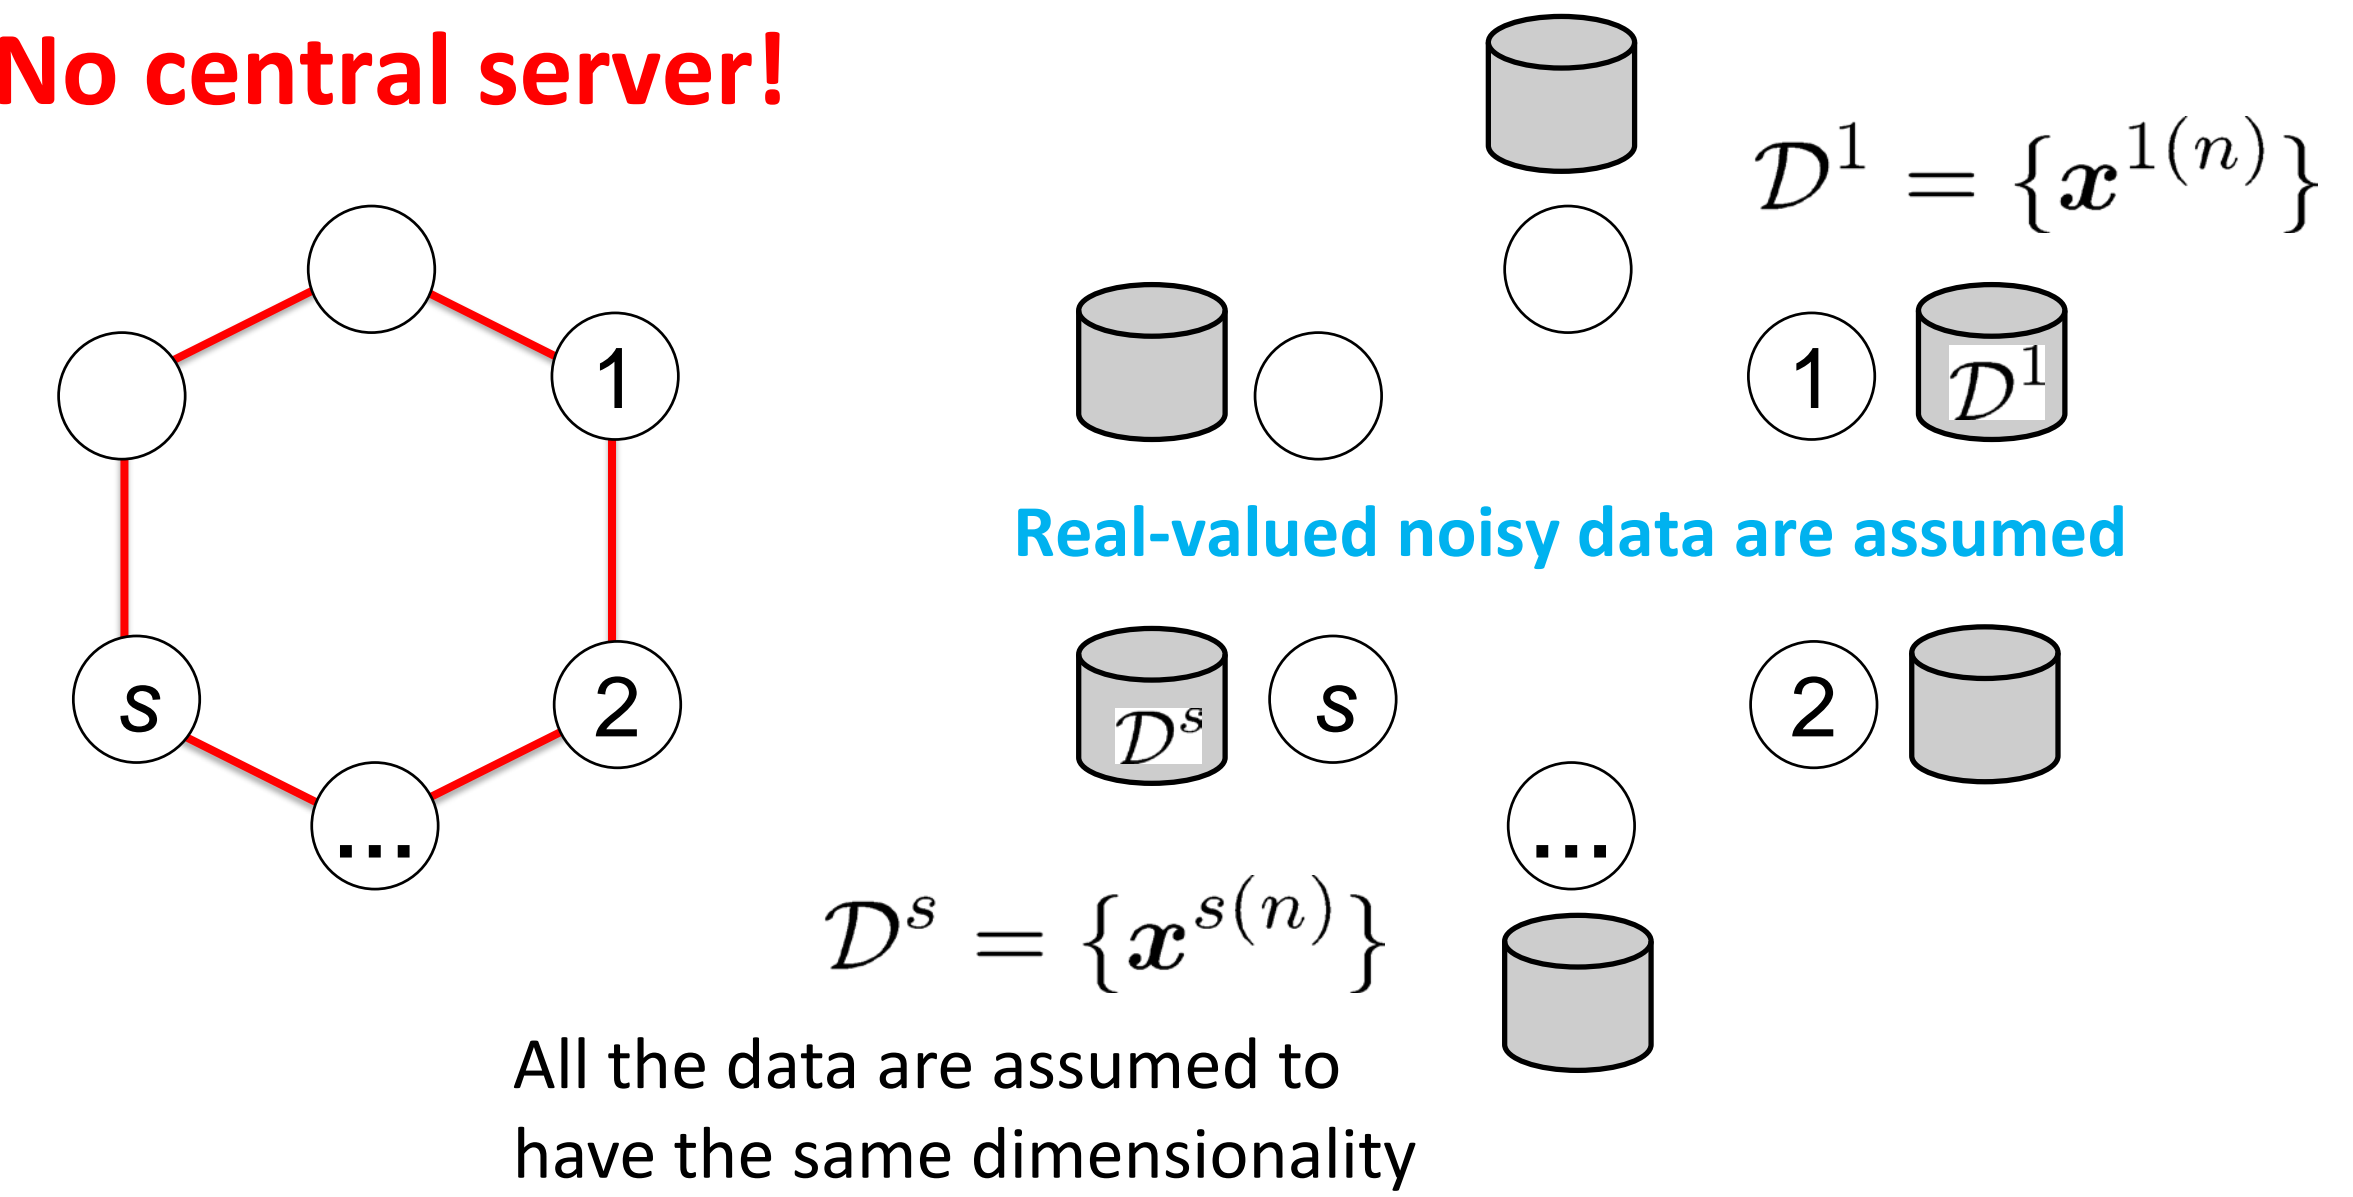


The problem was inspired by IBM Blockchain

Closed, membership-based, and decentralized network

Multiple "semi-honest" agents privately keep own data

No central server!



Prior work

Multi-task learning

- Actively studied area but mostly for supervised learning
- Not many of them are fully probabilistic
- Little is known about how to decentralize it

Decentralized computation

- Byzantine protocols assume categorical values
- Multi-agent consensus methods are not in the context of multi-task learning

Data privacy (under distributed environment)

- Differential privacy is problematic in distributed environment
- Secure multi-party computation typically needs a central server
- Homomorphic encryption is slow

Multi-task density estimation model

Employ a mixture model with agent-specific weights

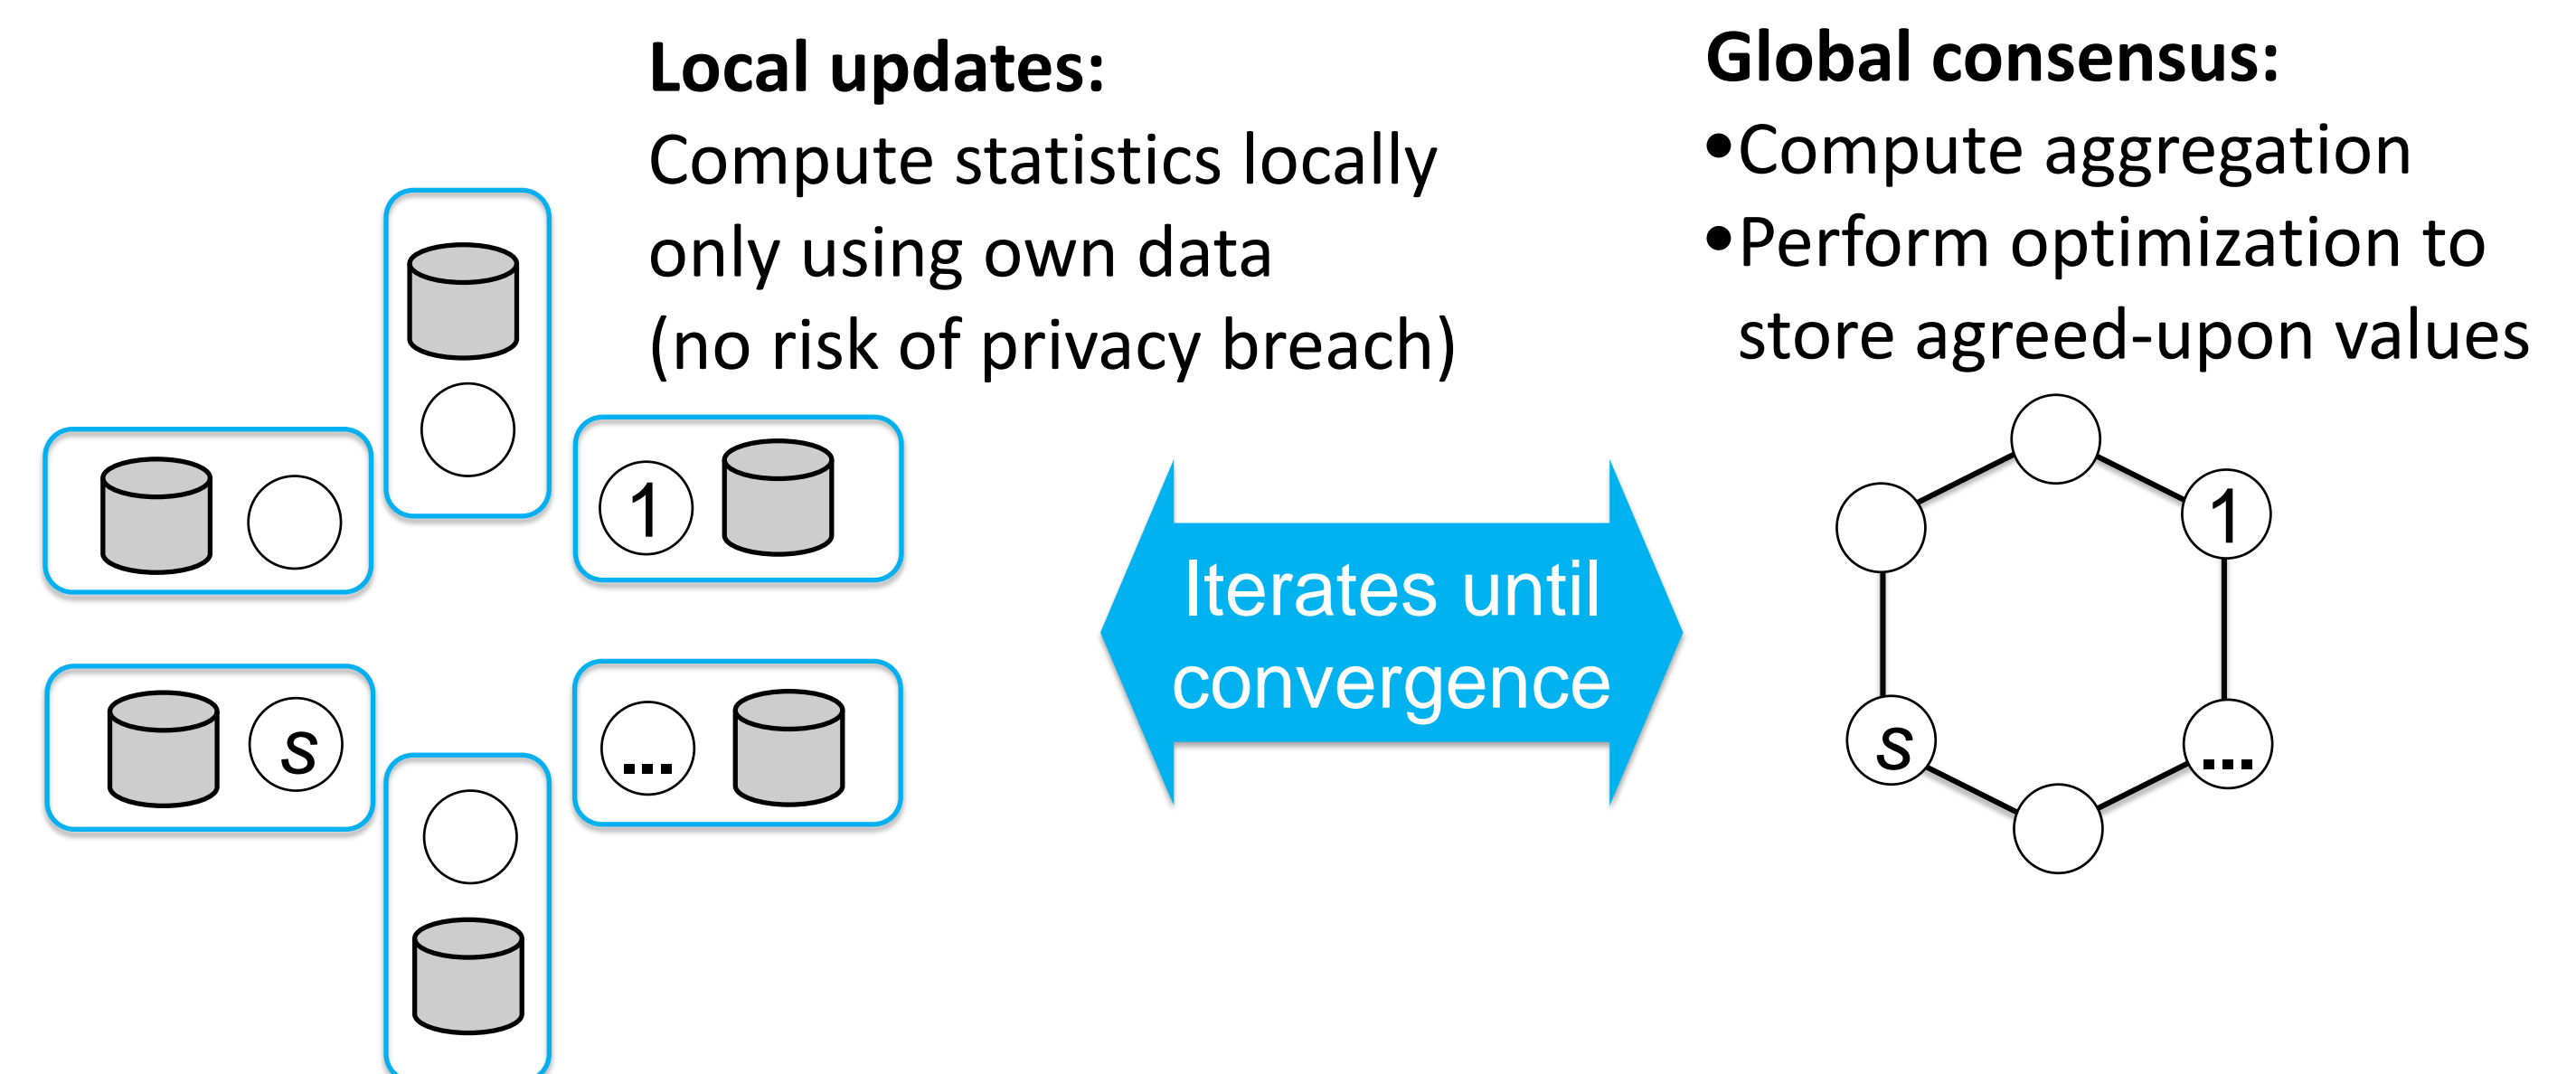
$$p^s(x | \Theta, \Pi^s) = \sum_{k=1}^K \pi_k^s f(x | \theta_k)$$

Shared by all the agents θ_k
Mixture weights are agent-specific π_k^s

The density is assumed to be in the exponential family

$$f(x | \theta_k) = G(\theta_k) H(x) \exp \{ \eta(\theta_k)^\top T(x) \}$$

Exponential family naturally leads to Global-Local Separation in maximum likelihood



Decentralized aggregation with data privacy

Decentralized aggregation
= Finding stationary state of Markovian transition process

Given the incidence matrix \mathbf{A} , an update equation

$$c^s \leftarrow c^s + \epsilon \sum_{j=1}^S A_{s,j} (c^j - c^s) \quad \text{or} \quad \mathbf{c} \leftarrow \mathbf{W}_\epsilon \mathbf{c}$$

converges to $\bar{\mathbf{c}} = \sum_{s=1}^S c^s = \mathbf{1}^\top \mathbf{c}$
S-dimensional vector of ones

$\mathbf{W}_\epsilon \equiv \mathbf{I} - \epsilon(\mathbf{D} - \mathbf{A})$
This is the graph Laplacian (minimum eigenvalue is 0)

"Chunking" method to prevent privacy breach

Randomly split each datum into N_c chunks and run aggregation algorithm N_c times (Simple!)

$$\bar{c} = \sum_{s=1}^S c^s$$

$c^s = c^{s[1]} + c^{s[3]} + c^{s[3]}$
 $\therefore \bar{c} = \bar{c}^{[1]} + \bar{c}^{[2]} + \bar{c}^{[3]}$

Orders of magnitude faster than homomorphic encryption-based methods

Note: Each chunking round has to use a different communication graph every chunking round.


- Can be done simply by shuffling network addresses (needs network router's help)
- Privacy breach probability can be made negligible by taking a large N_c

Graph structure matters!

What kind of communication graph \mathbf{A} should be chosen?


Cycle graph

- Most sparse and symmetric
- Slow convergence (quadratic in S)



"Cycle graph with inverse chord"

- Not regular/symmetric
- FAST convergence ($\log S$)



Motivating application:

collaborative condition-based monitoring of industrial assets

- IoT data is generally noisy
- Data privacy is a major concern but sample-wise encryption is not practical
- Need a new approach!

