Differentially Private Distributed Optimal Power Flow

Vladimir Dvorkin[1,2], Pascal Van Hentenryck[2], Jalal Kazempour[1], Pierre Pinson[1]

[1]Department of Electrical Engineering, Technical University of Denmark

[2] H. Milton Stewart School of Industrial and Systems Engineering (ISyE), Georgia Institute of Technology

Differential privacy

- Strong framework that ensures privacy of individuals when computing queries on datasets

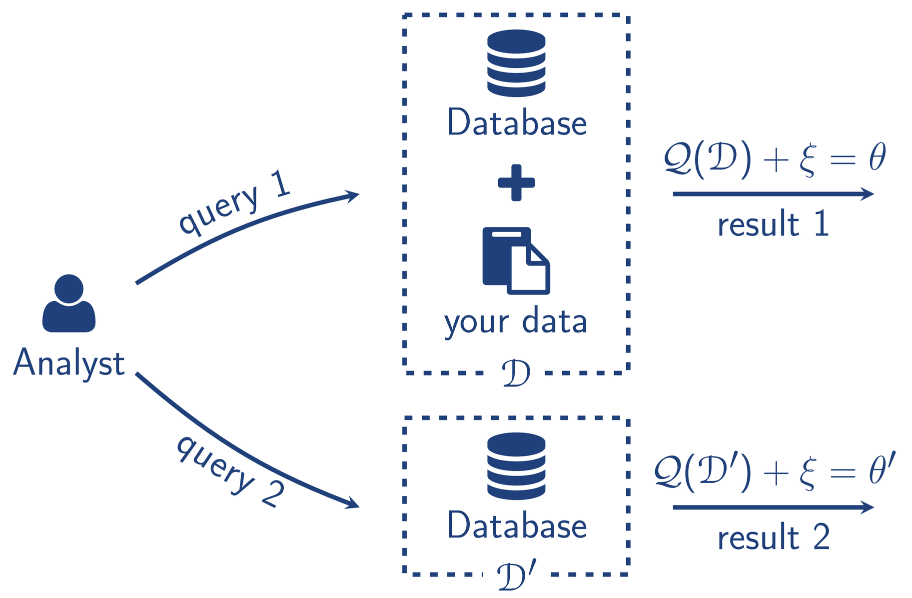

- $\mathcal{Q}$ is a query computed on a dataset
- ξ is a carefully calibrated noise
- θ and θ' are stat. indistinguishable
- By observing θ or θ', analyst can't tell if your data is included

ϵ-differential privacy

A randomized query $\tilde{\mathcal{Q}} : \mathcal{S} \mapsto \mathcal{R}$ with domain $\mathcal{S}$ and range $\mathcal{R}$ preserves $\epsilon-$differential privacy if for any output $\Theta \in \mathcal{R}$ and all adjacent datasets $\mathcal{D} \in \mathcal{S}$ and $\mathcal{D}' \in \mathcal{S}$, it holds that

$$\mathbb{P}[\tilde{\mathcal{Q}}(\mathcal{D}) \in \Theta] \leqslant \mathbb{P}[\tilde{\mathcal{Q}}(\mathcal{D}') \in \Theta]\exp(\epsilon),$$

where probability is taken over runs of $\tilde{\mathcal{Q}}$.

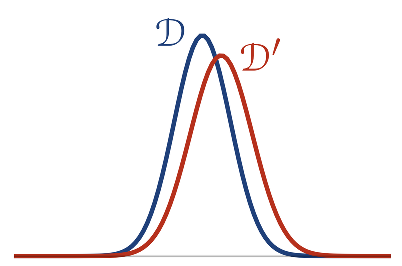

query outcomes

Distributed optimal power flow

- DC-OPF is distributed by duplicating voltage angles and dualizing consensus constraints

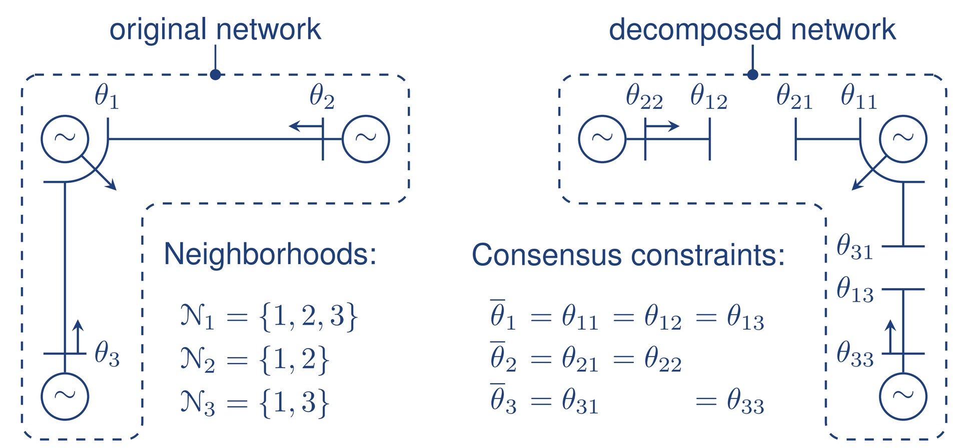

- Distributed ADMM-based optimal power flow algorithm writes as:

Agent (node) update $\quad \theta_i \leftarrow \underset{(p,\theta)\in\mathcal{O}_i}{\text{argmin}} \quad c_i(p_i) - \mu_i^\top \theta_i + \frac{1}{2}\|\overline{\theta} - \theta_i\|_\rho^2$

Consensus update $\quad \overline{\theta}_i \leftarrow \underset{\overline{\theta}_i}{\text{argmin}} \quad \mu_i^\top \overline{\theta}_i + \frac{1}{2}\|\overline{\theta} - \theta_i\|_\rho^2$

Dual update $\quad \mu_i \leftarrow \mu_i + \rho(\overline{\theta} - \theta_i)$

- The algorithms solely requires exchanging primal and dual variables across ADMM iterations
- **Does ADMM preserve privacy of local data, e.g., loads?**

Privacy concerns in distributed optimal power flow

- Although local data is not exchanged across ADMM iterations, under certain conditions it can be deduced from responses (θ_i) of agents to input signals ($\overline{\theta}_i; \mu_i$)
- From KKT conditions of agent subproblems, the load estimate $\hat{d}_i$ is obtained as

$$\hat{d}_i = [\mu_i + \rho(\overline{\theta}_j - \theta_i) - c_{1i}B_i][c_{2i}B_i]^{-1} - B_i^\top \theta_i,$$

provided that generator and transmission limits are not binding

- Hence, an adversary with side information (i.e., cost function, transmission data, and ADMM parameters) can deduce agent loads by intercepting agent communications
- Load inference on the IEEE 14-Bus Reliability Test System:

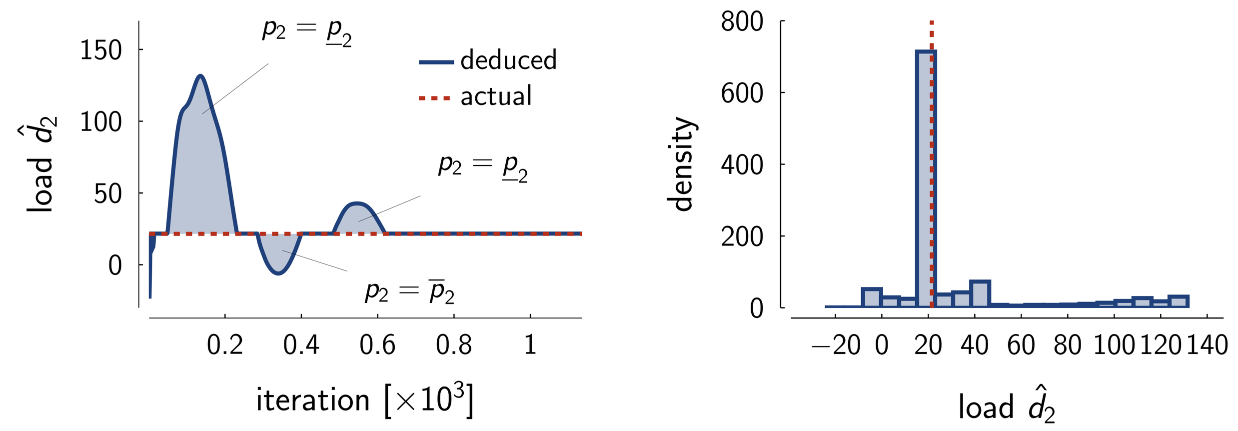

Differential privacy for distributed optimal power flow

To provide differential privacy, we treat agent optimization as query

$$\mathcal{Q}_i : \mathcal{D}_i \mapsto \theta_i,$$

that maps agent data into voltage angle estimates.

α-adjacent databases

We consider two agent databases $\mathcal{D}$ and $\mathcal{D}'$, different in load value by positive α, i.e.,

$$\|\mathcal{D}_i - \mathcal{D}_i'\|_1 \leqslant \alpha$$

ℓ_1-sensitivity of queries

The sensitivity of a query $\mathcal{Q}(\mathcal{D})$ amounts to

$$\Delta_{\mathcal{Q}} := \max_{\mathcal{D},\mathcal{D}'} \|\mathcal{Q}(\mathcal{D}) - \mathcal{Q}(\mathcal{D}')\|_1$$
$$\text{s.t. } \|\mathcal{D} - \mathcal{D}'\|_1 \leqslant \alpha,$$

Laplace mechanism

Let $\mathcal{Q} : \mathcal{D} \mapsto \mathbb{R}$ be a query that maps dataset $\mathcal{D}$ to reals. The randomized query

$$\tilde{\mathcal{Q}}(\mathcal{D}) = \mathcal{Q}(\mathcal{D}) + \xi$$

achieves $\epsilon-$differential privacy if ξ is sampled from the zero-mean Laplace distribution parametrized by privacy requirements and query sensitivity

$$\xi \sim \text{Lap}(\Delta_{\mathcal{Q}}/\epsilon)$$

Two deferentially private ADMM algorithms

Primal perturbation ADMM

$\min \mathcal{L}(p_i, \theta_i; \mu_i, \overline{\theta}_i) \rightarrow \theta_i + \xi_i = \tilde{\theta}_i$

$\min \mathcal{L}(\overline{\theta}; \mu, \tilde{\theta}) \rightarrow \overline{\theta}$

$\mu \leftarrow \mu + \rho(\overline{\theta} - \tilde{\theta})$

Dual perturbation ADMM

$\min \mathcal{L}(p_i, \theta_i; \mu_i + \xi_i, \overline{\theta}_i) \rightarrow \tilde{\theta}_i$

$\min \mathcal{L}(\overline{\theta}; \mu, \tilde{\theta}) \rightarrow \overline{\theta}$

$\mu \leftarrow \mu + \rho(\overline{\theta} - \tilde{\theta})$

- Both methods provably achieve differential privacy
- $\Delta_{\mathcal{Q}}$ is independent from $\overline{\theta}$ and μ
- No privacy loss accumulates across iterations
- OPF feasibility is not affected

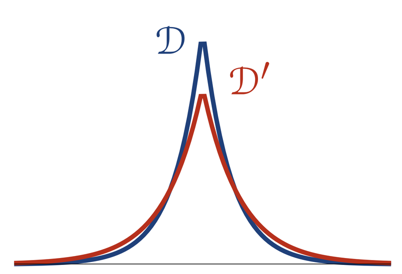

Numerical experiments

- We show the practical use of differential private ADMM algorithms in the face of adversarial load inference on a series of NESTA testbeds
- Load inference at selected nodes of IEEE 118-Bus Reliability Test System for varing adjacency coefficient α_i. By choosing α_i, the nodes decide to hide a certain portion of load. After running private ADMM 1000 times, we obtained the following distributions of the estimated loads:

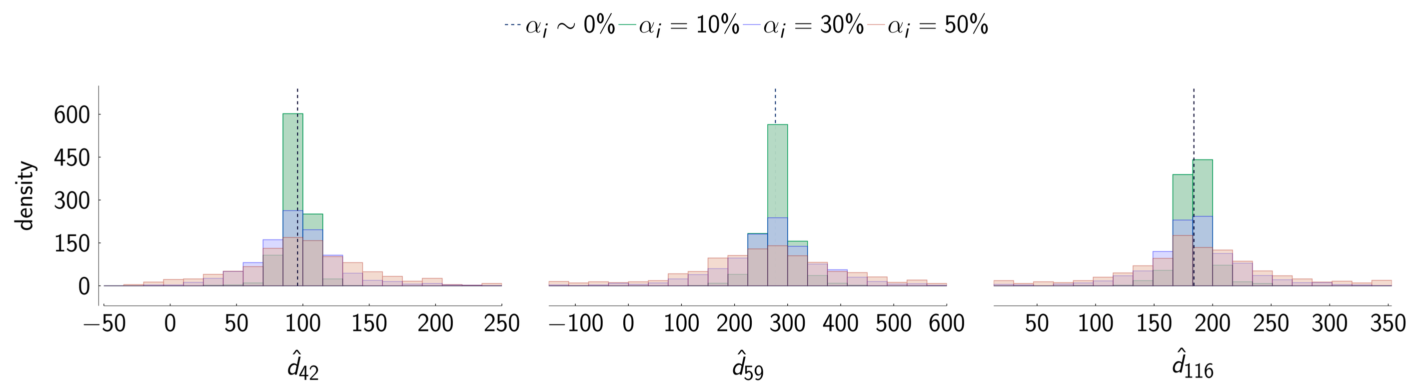

- Primal and dual variable perturbation methods achieve the same probability density of the output

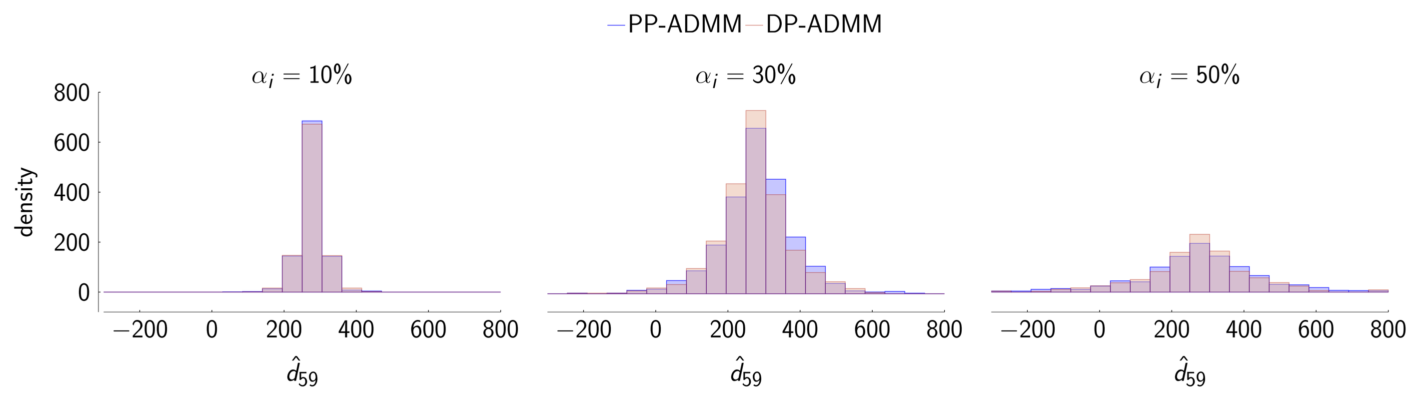

- Average optimality loss when providing privacy for PV-nodes for varying adjacency α. All in %:

case/α	5	10	15	case/α	5	10	15
3_lmbd	0.0	0.1	0.3	30_fsr	0.0	0.1	0.2
5_pjm	0.8	3.2	5.1	30_ieee	0.1	0.3	0.5
14_ieee	0.2	0.6	1.1	39_epri	0.1	0.4	1.1
24_ieee	2.0	5.3	9.0	57_ieee	0.2	2.2	3.5
30_as	1.5	5.1	5.4	118_ieee	3.4	8.1	11.4

- The system wins more privacy than loses optimality

Convergence statistics

case	-ADMM	$\alpha_i \sim 0\%$	$\alpha_i = 10\%$ min	avr	max
3_lmbd	PP	42	42	42	42
	DP		42	42	42
5_pjm	PP	85	75	95	175
	DP		75	84	116
14_ieee	PP	492	457	491	518
	DP		460	491	520
24_ieee	PP	771	316	735	1040
	DP		314	726	1430
30_as	PP	440	320	401	460
	DP		321	410	478
30_fsr	PP	247	247	247	247
	DP		247	247	247
30_ieee	PP	855	834	855	874
	DP		750	854	989
39_epri	PP	2320	1973	2307	2387
	DP		2237	2316	2370
57_ieee	PP	1679	1671	2050	2525
	DP		1545	2084	2658
118_ieee	PP	1836	1673	2007	2515
	DP		1596	3219	12526

- The two algorithms demonstrate similar convergence statistics
- Non-congested networks, e.g. 3_lmbd and 30_fsr, are immune to privacy preservation
- In the congested networks, the computational complexity remains the same only in expectation
- Both algorithms require extra amount of iterations compared to the privacy-agnostic ADMM.

Key messages

- ADMM *fails* to preserve privacy unless augmented by privacy-cognizant practices
- We introduce two differentially private ADMM for distributed OPF that keeps data on loads private in the face of adversarial inference from communication signals
- The key idea is to apply a calibrated noise to agent responses to hide items in agent datasets:
  - The noise does not affect the OPF feasibility
  - Convergence is achieved on all test cases
  - More insights on convergence properties are required
- There exists privacy-optimality trade-offs that raise institutional issues (e.g., price of privacy)

References

Dvorkin, V., Van Hentenryck, P., Kazempour, J., & Pinson, P. (2019). Differentially Private Distributed Optimal Power Flow. arXiv preprint arXiv:1910.10136.