Synthesizing Grid Data with Cyber Resilience and Privacy Guarantees

Shengyang Wu, Vladimir Dvorkin

Department of Electrical Engineering and Computer Science

University of Michigan

2025 INFORMS Annual Meeting — Atlanta

October 26, 2025

UNIVERSITY OF MICHIGAN

Power systems as a stroll in the fog

- Power systems are critical infrastructures with most of data being classified
- We have only a limited observably, e.g., ISO data disclosure portals
- Grid stakeholders hence act on a limited set of system data

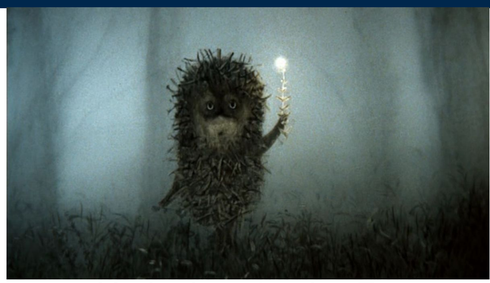

Hedgehog in the Fog
Yuri Norstein (1975)

Example: DC optimal power flow (OPF) in the (small) IEEE 118-bus system

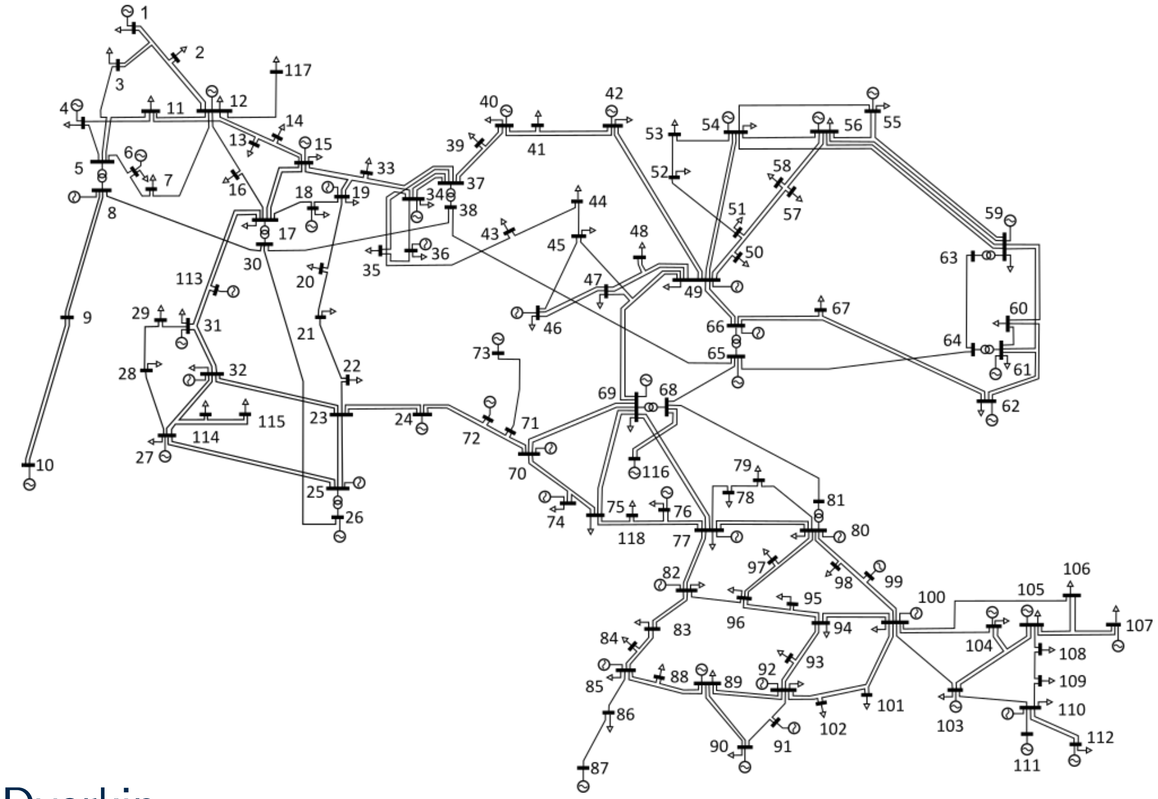

- 1079 rows of element-specific data
- Each generator owns only 2 rows
- The rest of system data is not *explicitly* disclosed to power producers

How to disclose grid data in a controllable manner?

Differential privacy (DP) enables controlled disclosure of grid data

Differential Privacy for Power Grid Obfuscation

Ferdinando Fioretto, Terrence W. K. Mak, *Member, IEEE*, and Pascal Van Hentenryck, *Member, IEEE*

Abstract—The availability of high-fidelity energy networks bring significant value to academic and commercial research. However, such releases also raise fundamental concerns related to privacy and security as they can reveal sensitive commercial information and expose system vulnerabilities. This paper investigates how to release the data for power networks where *the parameters of transmission lines and transformers are obfuscated*, it does so by using the framework of Differential Privacy (DP), that provides strong privacy guarantees and has attracted significant attention in recent years. Unfortunately, simple DP mechanisms often result in AC-infeasible networks. To address these concerns, this paper presents a novel differentially private mechanism that guarantees AC-feasibility and largely preserves the fidelity of the obfuscated power network. Experimental results also show that the obfuscation significantly reduces the potential damage of an attack carried by exploiting the released dataset.

Index Terms—Data privacy, power system security, data processing.

I. Introduction

THE AVAILABILITY of test cases representing high-fidelity power system networks is fundamental to foster research in optimal power flow, unit commitment, and transmission planning, to name only a few challenging problems. This need was recognized by ARPA-E when it initiated the *Grid Data Program* in 2015. However, the release of such rich datasets is challenging due to legal issues related to privacy and national security. For instance, the electrical load of an industrial customer indirectly exposes sensitive information on its production levels and strategic investments, and the value parameters of lines and generators may reveal how transmission operators operate their networks. Furthermore, this data could be exploited by an attacker to inflict targeted damages on the network infrastructure.

This paper explores whether differential privacy (DP) [1] is an algorithmic property that measures and bounds the privacy risks associated with answering sensitive queries or releasing a privacy-preserving dataset. It introduces carefully calibrated noise to the data to prevent the disclosure of sensitive information. An algorithm satisfying DP offers privacy protection regardless of the external knowledge of an attacker. In particular, the definition of DP adopted in this paper ensures that an attacker obtaining access to a differentially private output, cannot detect (with high probability) how close is the privacy-preserving value to its original one.

However, DP faces significant challenges when the resulting privacy-preserving datasets are used as inputs to complex optimization algorithms, e.g., *Optimal Power Flow* (OPF) problems. Indeed, the privacy-preserving dataset may have lost the fidelity and realism of the original data and may even not admit feasible solutions for the optimization problems of interest [2].

This paper studies how to address such a challenge when the goal is to preserve the privacy of line parameters and transformers. It presents a DP mechanism to release data for power networks that is realistic and limits the power of an attacker. More precisely, the contribution of this paper is fourfold. *(1)* It proposes the *Power Line Obfuscation* (PLO) mechanism to obfuscate the line parameters in a power system network. *(2)* It shows that PLO has strong theoretical properties: It achieves ϵ-differential privacy, it ensures that the released data can produce feasible solutions for OPF problems, and its objective value is a constant factor away from optimality. *(3)* It extends the PLO mechanism to handle time-series network data. *(4)* It demonstrates, experimentally, that the PLO mechanism improves the accuracy of existing approaches. When tested on the largest collection of OPF test cases available, it results in solutions with similar costs and optimality gaps to those obtained on the original problems, while also protecting well against an attacker that has access to the released data and uses it to damage the real power network. Interestingly, on the test cases adopted, the damage inflicted on the real power

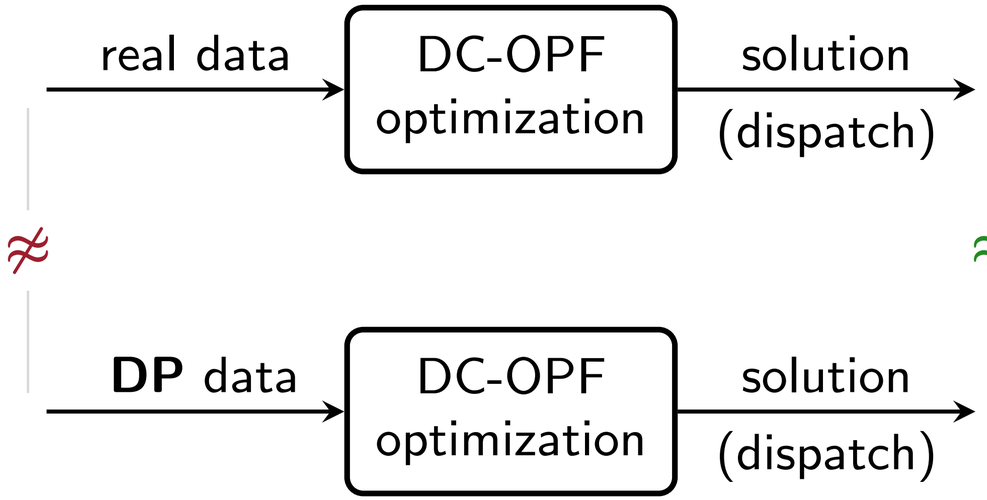

IEEE CSS IEEE CONTROL SYSTEMS LETTERS, VOL. 7, 2023

Differentially Private Algorithms for Synthetic Power System Datasets

Vladimir Dvorkin Jr., *Member, IEEE*, and Audun Botterud, *Member, IEEE*

Abstract—While power systems research relies on the availability of real-world network datasets, data owners (e.g., system operators) are hesitant to share data due to privacy risks. To control these risks, we develop privacy-preserving algorithms for synthetic generation of datasets for optimization and machine learning. Taking a real-world dataset as input, the algorithms output its noisy, synthetic version, which preserves the accuracy of the real data on a specific downstream model or even a large population of those. We control the privacy loss using Laplace and Exponential mechanisms of differential privacy and preserve data accuracy using a post-processing convex (or mixed-integer) optimization. We apply the algorithms to generate synthetic power network parameters and wind power data.

Index Terms—Differential privacy, machine learning, power systems optimization, synthetic datasets.

I. Introduction

POWER system datasets are instrumental for enhancing solutions to many problems, including optimal power flow (OPF) and wind power forecasting. Releasing real data, however, is challenging due to security and privacy concerns. Indeed, detailed network datasets inform cyberattacks on SCADA systems and can be used by strategic market players to maximize profits at the expense of deteriorating social welfare. These concerns motivate producing synthetic datasets – a sanitized version of private datasets that approximately preserve accuracy of data for power system applications.

Differential privacy (DP) is an algorithmic notion of privacy preservation that enables trade-offs between data privacy and accuracy in optimization [1] and machine learning [2]. It also found applications in the context of privacy-preserving OPF computations, e.g., in distributed algorithms [3] and centralized solvers for power grids [4], [5], as well as in machine learning problems specific to power systems [6]. Models in [3], [4], [5], [6], however, only control data leakages in computations and do not provide synthetic datasets per se.

Producing synthetic datasets in a DP way is achieved by corrupting data with privacy-preserving noise [7], [8]. However, applications of the standard noise-additive DP mechanisms in power systems, such as the Laplace mechanism, may no longer admit a meaningful result. Indeed, adding noise to data may fundamentally alter important statistics and trends in machine learning datasets [9]. In the OPF context, [10] and [11] showed that the Laplacian perturbation of network data almost surely violates OPF feasibility. As a remedy, they proposed an optimization-based post-processing which restores the accuracy of synthetic OPF datasets without altering the privacy guarantee. The proposed restoration, however, renders the Laplace mechanism feasible only for a particular instances of variable renewable production) may not be possible, as noise must be scaled drastically as per composition of DP [12].

In this letter, we introduce private synthetic dataset generation algorithms for power systems, which ensure the accuracy of synthetic datasets for downstream models. They enjoy known DP mechanisms and convex (or mixed-integer) post-processing optimization of data. Specifically, we develop:

1) Wind Power Obfuscation (WPO) algorithm, which privately generates wind power measurements, while guaranteeing DP of the real data and ensuring accuracy in terms of the outcomes of a regression analysis.
2) Transmission Capacity Obfuscation (TCO) algorithm, which generates synthetic line parameters, while ensuring their feasibility and cost-consistency on a population of OPF models. Here, we use both Laplace and Exponential mechanisms of DP to substantially reduce the noise compared to using the Laplace mechanism alone.

- DP principle: obfuscate real data (add noise) but preserve its value
- In the DC-OPF setting: obfuscate grid data but preserve the OPF solution
- Formal *privacy guarantee*: released DP data does NOT disclose the real data
- Many applications to synthesizing high-quality transmission, load and generation data

Example: Synthesizing transmission line capacities using DP and optimization

- IEEE 73-RTS benchmark
- **Step 1**: add random noise to trans capacity

$$\varphi_1 = \overline{f} + \text{calibrated noise}$$

- **Step 2**: post-process φ_1 to ensure OPF feasibility and cost-consistency w.r.t. real trans capacity

$$\begin{aligned} \varphi_2 \in \operatorname*{argmin}_{\varphi} \quad & \|c(\overline{f}) - c(\varphi)\| + \|\varphi - \varphi_1\| \\ \text{s.t.} \quad & c(\varphi) = \min_{p} \quad c(p) \quad \textit{opf cost} \\ \text{s.t.} \quad & p \in \mathcal{P}(\varphi) \quad \textit{opf feas} \end{aligned}$$

- **Step 3-N**: repeat Step 2 until OPF feasibility and cost-consistency are restored across many scenarios

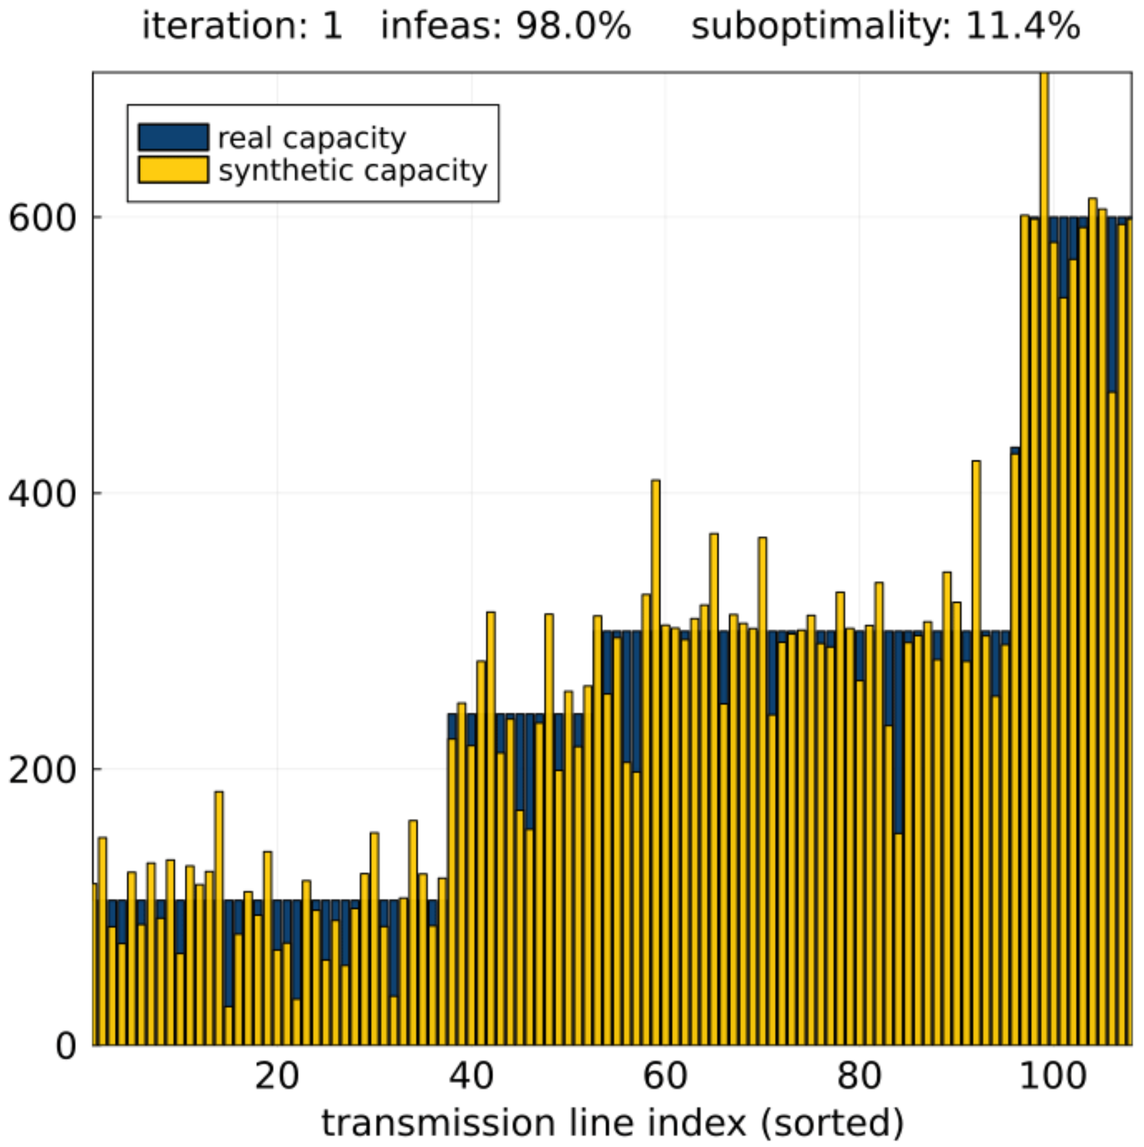

Dvorkin, V., Botterud A. Differentially private algorithms for synthetic power system datasets, IEEE Control Systems Letters, 2023.

Challenge: If DP synthetic datasets are so good, do they pose any security risks?

- Grid datasets are used for calibrating cyberattacks on power grids
- Hypothesis: high-quality synthetic data $\rightarrow$ well-calibrated attacks
- Classes of cyberattacks: false data injection, line outage masking, physical attacks, ...

Contribution: We identify cyberattack risks in releasing DP grid data and propose new algorithms to guarantee both privacy and cyber resilience to source grids

Shengyang Wu

Load redistribution attack

- Given load $\mathbf{d}$, the attack corrupts the load $\mathbf{d} + \boldsymbol{\delta}$ with bus injection $\boldsymbol{\delta}$ from admissible set Δ

$$\Delta \triangleq \left\{ \boldsymbol{\delta} \;\middle|\; \begin{array}{ll} \underline{\boldsymbol{\delta}} \leqslant \boldsymbol{\delta} \leqslant \overline{\boldsymbol{\delta}} & \textit{injection limits for each bus} \\ \mathbf{1}^\top \boldsymbol{\delta} = 0 & \textit{total load remains unchanged} \end{array} \right\}$$

- Amounts solving a bi-level optimization problem:

$$\begin{aligned} C_{\text{att}}^{\text{BO}}(\mathbf{d}) = \max_{\boldsymbol{\delta} \in \Delta} \quad & C_{\text{opf}}(\mathbf{d} + \boldsymbol{\delta}) && \textit{maximize the cost} \\ \text{s.t.} \quad & C_{\text{opf}}(\mathbf{d} + \boldsymbol{\delta}) = \min_{\mathbf{x}} \quad \mathbf{c}^\top \mathbf{x} && \textit{feedback from OPF} \\ & \qquad\qquad \text{s.t.} \quad \mathbf{a}_k^\top \mathbf{x} + \mathbf{b}_k^\top (\mathbf{d} + \boldsymbol{\delta}) + e_k \leqslant \mathbf{0} \end{aligned}$$

- The problem seeks a *stealthy* attack vector that maximizes the OPF cost

Can DP grid data be used to successfully execute the load redistribution attack?

Numerical evidence

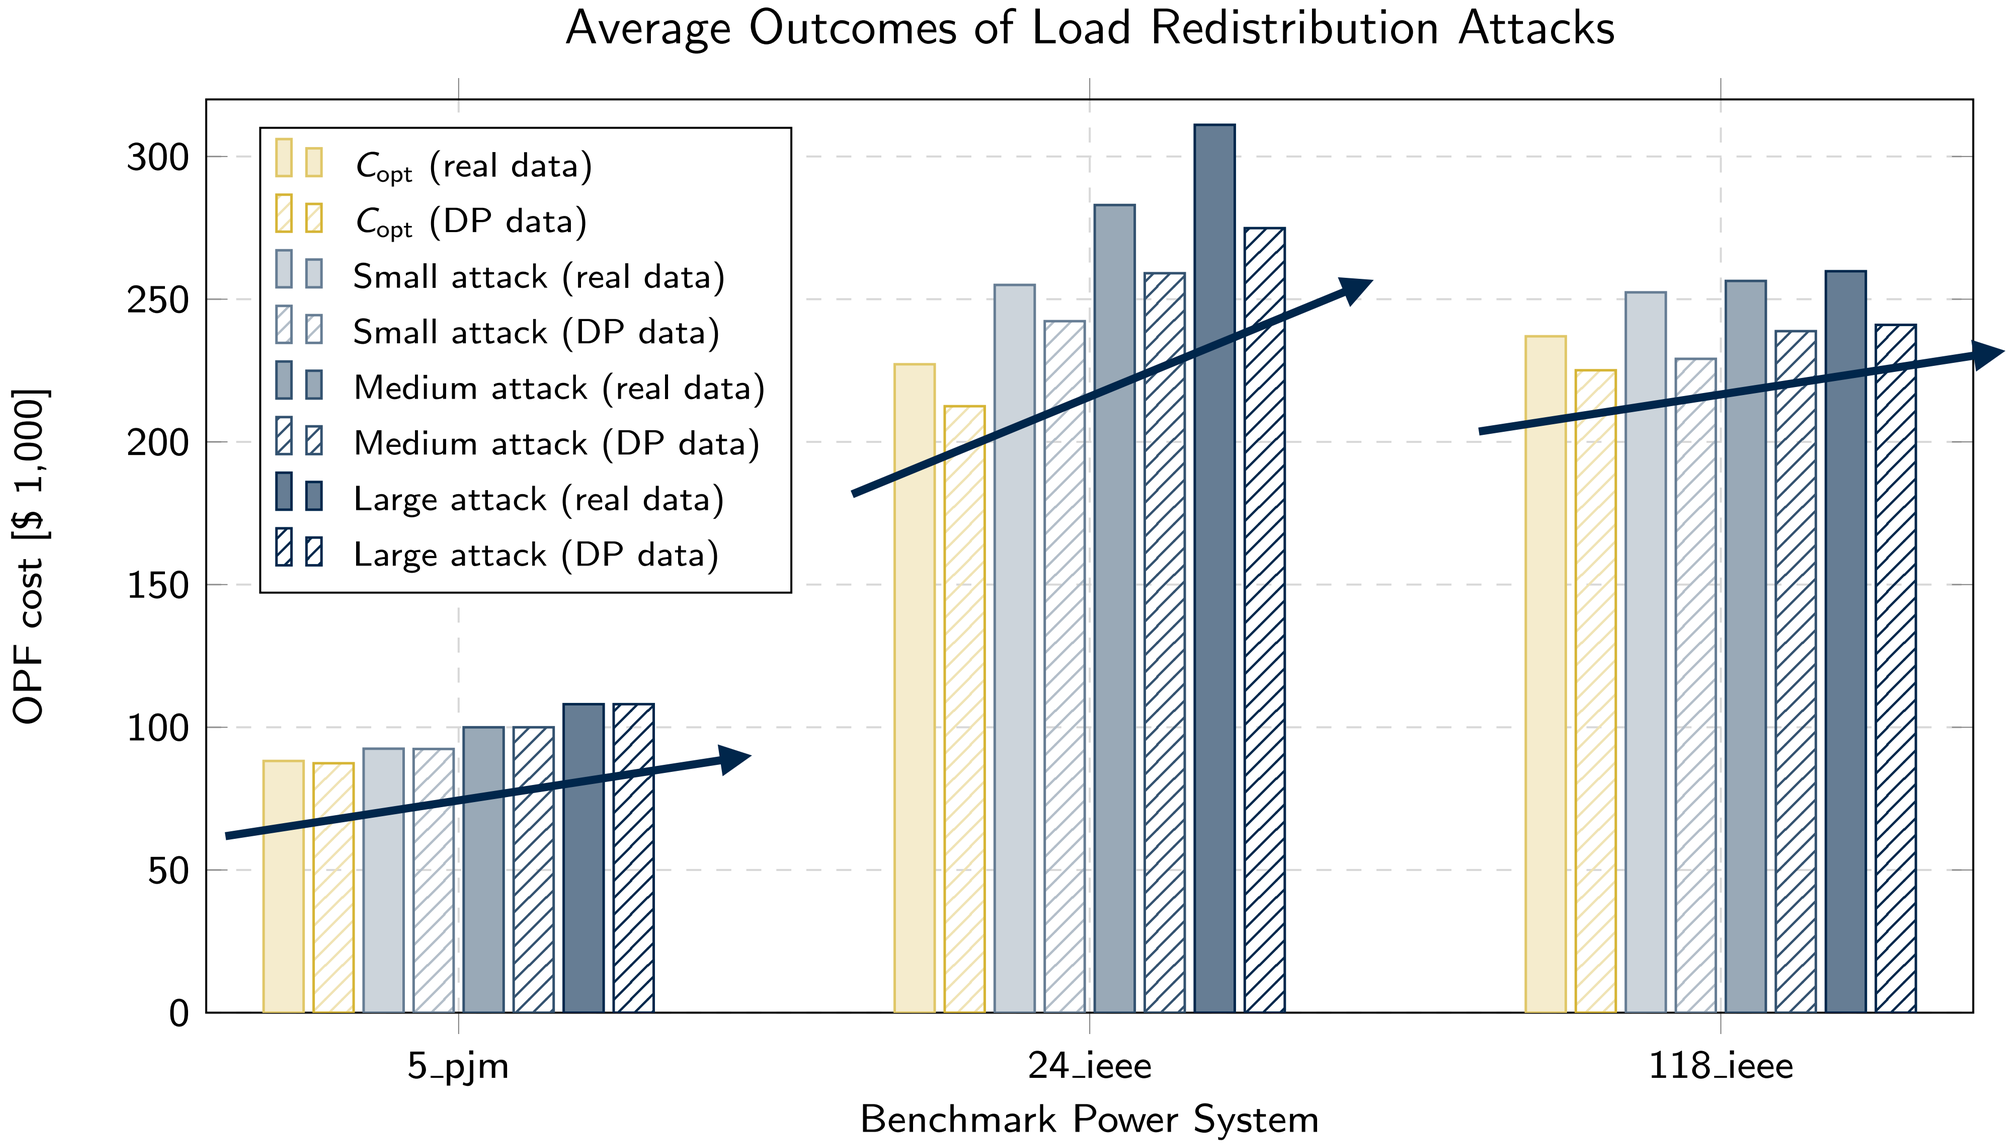

Post-processing optimization to minimize attack damage

$$
\begin{aligned}
&\underset{\mathbf{d}}{\text{minimize}} && \underbrace{\|C_{\text{att}}^{\text{BO}}(\mathbf{d}) - \tilde{C}_{\text{opf}}\|}_{\text{power of attack}} + \beta \underbrace{\|C_{\text{opf}}(\mathbf{d}) - \tilde{C}_{\text{opf}}\|}_{\text{cost consistency}} + \gamma \underbrace{\|\mathbf{d} - \mathbf{d}_1\|}_{\text{regularization}} && \textbf{level 1}\\
&\text{subject to} && C_{\text{opf}}(\mathbf{d}) = \\
&&& \quad \underset{x}{\text{minimize}} \quad C_{\text{opf}}(x) && \textbf{level 2}\\
&&& \quad \text{subject to} \quad x \in \text{opf-eq}(\mathbf{d}) \\
&&& C_{\text{att}}^{\text{BO}}(\mathbf{d}) = \\
&&& \quad \underset{\delta \in \Delta}{\text{maximize}} \quad C_{\text{opf}}(\delta) \\
&&& \quad \text{subject to} \quad C_{\text{opf}}(\delta) = \\
&&& \qquad \underset{x}{\text{minimize}} \quad C_{\text{opf}}(x) && \textbf{level 3}\\
&&& \qquad \text{subject to} \quad x \in \text{opf-eq}(\mathbf{d} + \delta)
\end{aligned}
$$

Figure: Tri-level structure of the cyber-resilient post-processing of DP load data.

- **Step 1**: add random noise to real loads: $\mathbf{d}_1 = \mathbf{d} +$ calibrated noise
- **Step 2**: post-process $\mathbf{d}_1$ by solving a tri-level optimization:
  - Level-1: Optimize synthetic load $\mathbf{d}$ to balance attack damage and cost-consistency
  - Level-2: Feedback from both OPF and attack optimization
  - Level-3: Embedded OPF for attack calibration
- **Result**: DP load vector $\mathbf{d}$ balancing attack damage and cost-consistency

Tractable approximation of the tri-level problem

- Optimizing synthetic loads over bi-level optimization is computationally challenging
- We drawn on the connection between bi-level and robust (single-level) optimization

Bi-level attack optimization

$$C_{\text{att}}^{\text{BO}}(\mathbf{d}) = \max_{\boldsymbol{\delta} \in \Delta} C_{\text{opf}}(\mathbf{d} + \boldsymbol{\delta})$$
$$\mathbf{x} \in \operatorname*{argmin}_{\mathbf{x}} \ \mathbf{c}^\top \mathbf{x}$$
$$\text{s.t. } \mathbf{a}_k^\top \mathbf{x} + \mathbf{b}_k^\top (\mathbf{d} + \boldsymbol{\delta}) + e_k \leqslant \mathbf{0} \quad \forall k$$

(uniform attack)

Robust optimization (RO) approximation

$$C_{\text{att}}^{\text{RO}}(\mathbf{d}) = \min_{\mathbf{x}} \ \mathbf{c}^\top \mathbf{x}$$
$$\text{s.t.} \quad \max_{\boldsymbol{\delta}_k \in \Delta} \left[\mathbf{a}_k^\top \mathbf{x} + \mathbf{b}_k^\top (\mathbf{d} + \boldsymbol{\delta}_k) + e_k \right] \leqslant \mathbf{0} \quad \forall k$$

(per constraint attack)

Proposition: For any feasible load $\mathbf{d}$, relation $C_{\text{att}}^{\text{RO}}(\mathbf{d}) \geqslant C_{\text{att}}^{\text{BO}}(\mathbf{d})$ holds.

Cyber Resilient Obfuscation (CRO) Algorithm

- **Step 1**: Obfuscate real load data

  DP load $\mathbf{d}_1 = \mathbf{d} +$ calibrated noise

  DP estimation of OPF costs: $\tilde{C}_{\text{opf}} = C_{\text{opf}}(\mathbf{d}) +$ calibrated noise

- **Step 2**: Post-process $\mathbf{d}_1$ to balance cost-consistency and cyber resilience P

$$\underset{\mathbf{d}}{\text{minimize}} \quad \underbrace{\|C^{\text{RO}}_{\text{att}}(\mathbf{d}) - \tilde{C}_{\text{opf}}\|}_{\text{power of attack}} + \beta \underbrace{\|C_{\text{opf}}(\mathbf{d}) - \tilde{C}_{\text{opf}}\|}_{\text{cost consistency}} + \underbrace{\gamma\|\mathbf{d} - \mathbf{d}_1\|}_{\text{regularization}}$$

embedded attack optimization — **embedded OPF optimization**

- Replacing the two embedded optimization problems with their Karush–Kuhn–Tucker conditions leads to a single-level mixed-integer problem.

CRO-Exp: selecting only important constraints for post-processing

$$C^{\text{RO}}_{\text{att},\tau}(\mathbf{d}) = \min_{\mathbf{x}} \ \mathbf{c}^\top \mathbf{x}$$

$$\text{s.t.} \quad \max_{\delta_k \in \Delta} \left[\mathbf{a}_k^\top \mathbf{x} + \mathbf{b}_k^\top (\mathbf{d} + \delta_k) + e_k \right] \leqslant \mathbf{0} \quad \forall k \in \mathcal{K}' \qquad \textit{RO-reformulated cons}$$

$$\left[\mathbf{a}_k^\top \mathbf{x} + \mathbf{b}_k^\top \ \mathbf{d} \qquad + e_k \right] \leqslant \mathbf{0} \quad \forall k \in \mathcal{K} \qquad \textit{original problem cons}$$

- We select only most important τ constraints for RO reformulation
- Most important constraints in $\mathcal{K}'$ are function of original loads
- Exponential mechanism of DP to obfuscate loads when selecting $\mathcal{K}'$

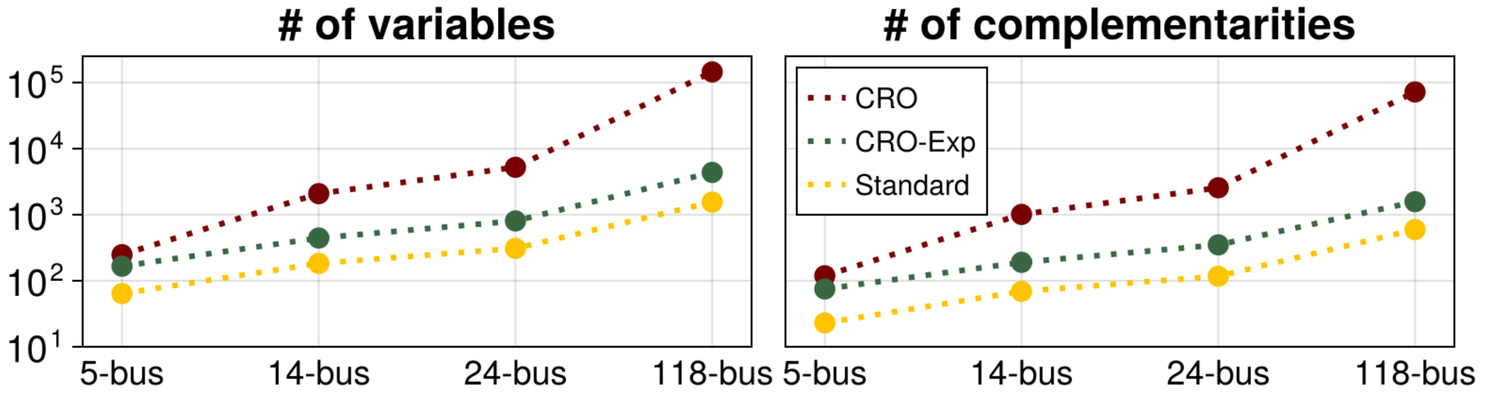

Selecting only important constraints substantially reduces the computational burden

Algorithm 1: Privacy-preserving CRO-Exp

Input: real load $\mathbf{d}$, DP parameters $(\alpha, \varepsilon_1, \varepsilon_2, \varepsilon_3)$, attack data $(\beta, \gamma, \Delta, \tau)$, $\mathcal{K} = \{\emptyset\}$

1 DP obfuscation of load and OPF costs:

$$\tilde{\mathbf{d}}^0 = \mathbf{d} + \text{Lap}\left(\frac{\alpha}{\varepsilon_1}\right)^n$$

$$\tilde{C}_{\text{opf}} = C_{\text{opf}}(\mathbf{d}) + \text{Lap}\left(\frac{\alpha\overline{c}}{\varepsilon_2}\right)$$

2 DP estimation of the set $\mathcal{K}$ of the worst-case constraints

for $t = 1, \ldots, \tau$ **do**
  for $k = 1, \ldots, K$ **do**
    $C_k = C^{\text{RO}}_{\text{att},t}(\mathbf{d}) + \text{Lap}\left(\frac{\alpha\overline{c}}{\varepsilon_3}\right)$
  end
  $k_t \leftarrow \text{argmax}_k \; C_k$
  $\mathcal{K} \leftarrow \mathcal{K} \cup \{k_t\}$
end

3 Post-processing optimization of the synthetic load vector

$$\tilde{\mathbf{d}} \in \underset{\tilde{\mathbf{d}}}{\text{argmin}} \; \|C_{\text{opf}}(\tilde{\mathbf{d}}) - \tilde{C}_{\text{opf}}\|_1 + \beta\|C^{\text{RO}}_{\text{att},\tau}(\tilde{\mathbf{d}}) - \tilde{C}_{\text{opf}}\|_1 + \gamma\|\tilde{\mathbf{d}} - \tilde{\mathbf{d}}^0\|_1$$

Output: Synthetic load vector $\tilde{\mathbf{d}}$

CRO application to PJM 5-Bus system

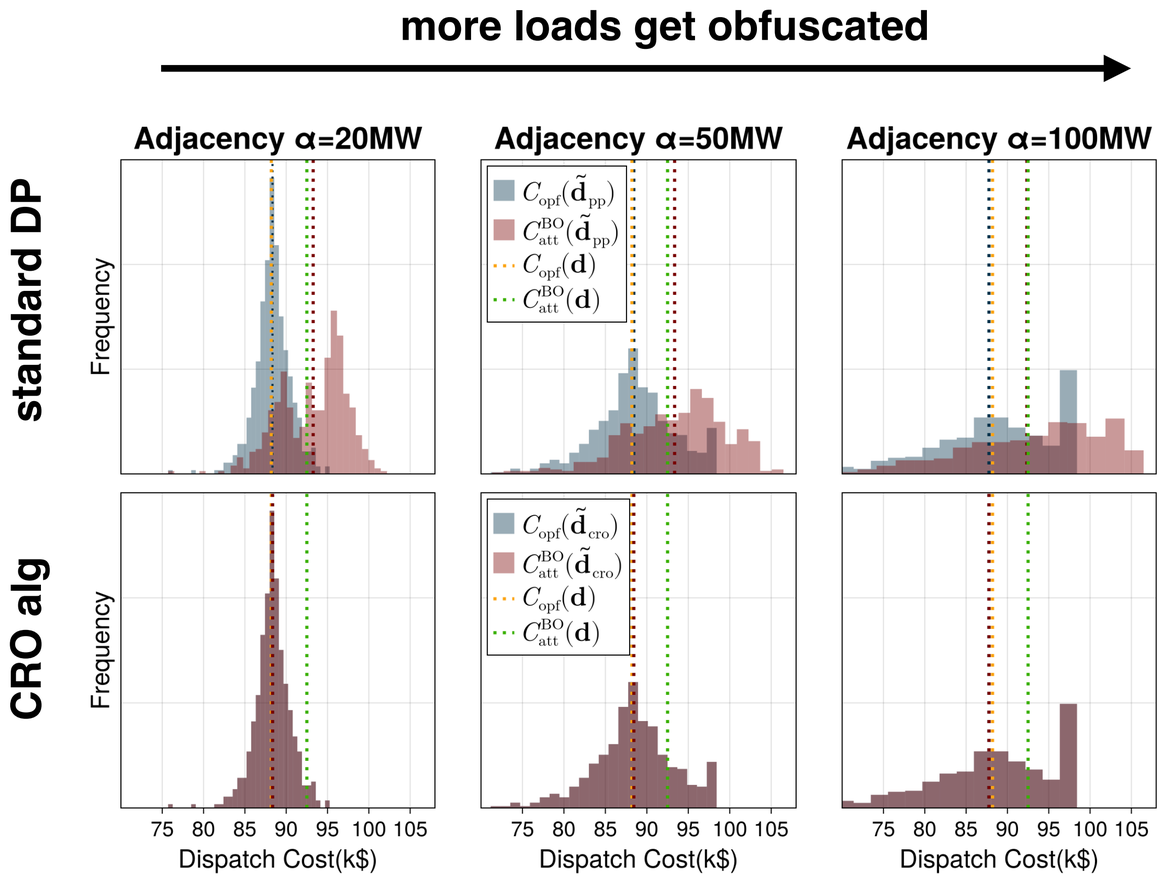

OPF cost distributions in normal and post-attack operation

- Blue distributions - normal operation
- Red distributions - post-attack operation
- Top row - standard DP post-processing
- Bottom row - proposed CRO post-processing

CRO sends important signal to attackers: **the attacks do not lead to extra OPF cost to the system**.
(normal and post-attack distributions overlap)

CRO-Exp application to larger systems

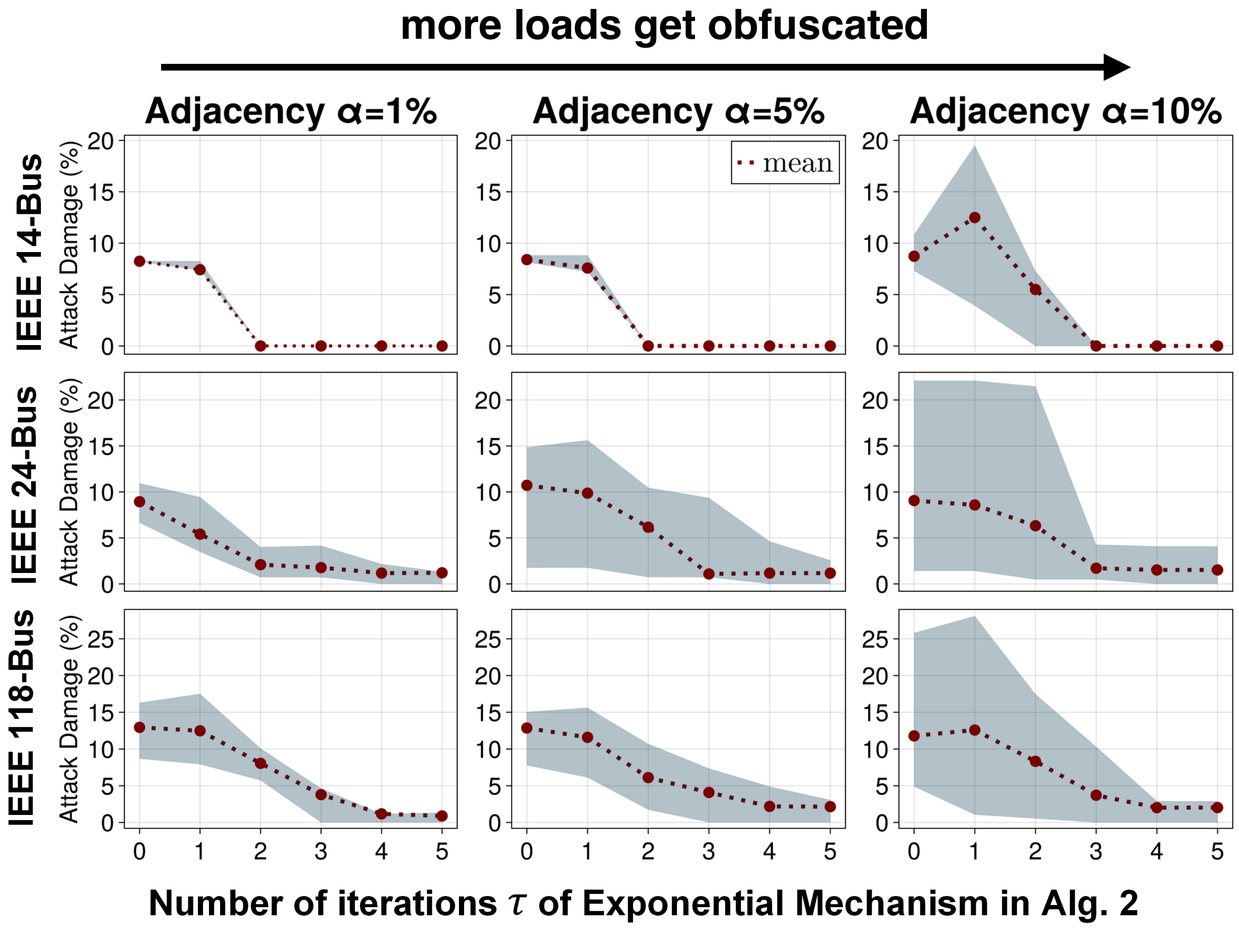

Post-attack damage as a function of constraints selected for post-processing optimization

- The more constraints under attack $\rightarrow$ more resilent the source grid to load redistribution attacks
- 5 constraints on average to minimize the damage

After 5 iterations, the CRO-Exp algorithm identified the most important constraints for post-processing synthetic loads and **ensuring grid resilience**.

Conclusions

- Synthetic grid data is optimized to guarantee privacy, quality and cyber resilience simultaneously
- Trade-offs under linear cost functions are "flat": resilience is achieved with little to no impact on data quality
- The tri-level post-processing optimization can be efficiently collapsed to single-level optimization under reasonable and judicious approximations (connecting bilevel and robust optimization techniques)

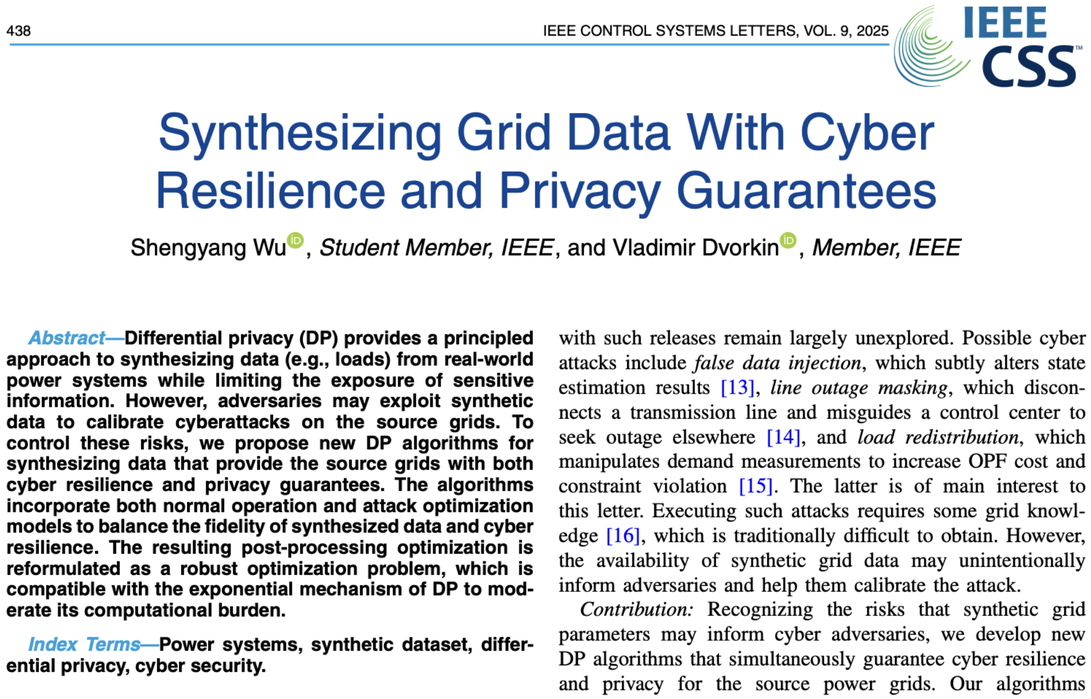

IEEE CSS

Synthesizing Grid Data With Cyber Resilience and Privacy Guarantees

Shengyang Wu, *Student Member, IEEE*, and Vladimir Dvorkin, *Member, IEEE*

***Abstract*—Differential privacy (DP) provides a principled approach to synthesizing data (e.g., loads) from real-world power systems while limiting the exposure of sensitive information. However, adversaries may exploit synthetic data to calibrate cyberattacks on the source grids. To control these risks, we propose new DP algorithms for synthesizing data that provide the source grids with both cyber resilience and privacy guarantees. The algorithms incorporate both normal operation and attack optimization models to balance the fidelity of synthesized data and cyber resilience. The resulting post-processing optimization is reformulated as a robust optimization problem, which is compatible with the exponential mechanism of DP to moderate its computational burden.**

***Index Terms*—Power systems, synthetic dataset, differential privacy, cyber security.**

with such releases remain largely unexplored. Possible cyber attacks include *false data injection*, which subtly alters state estimation results [13], *line outage masking*, which disconnects a transmission line and misguides a control center to seek outage elsewhere [14], and *load redistribution*, which manipulates demand measurements to increase OPF cost and constraint violation [15]. The latter is of main interest to this letter. Executing such attacks requires some grid knowledge [16], which is traditionally difficult to obtain. However, the availability of synthetic grid data may unintentionally inform adversaries and help them calibrate the attack.

Contribution: Recognizing the risks that synthetic grid parameters may inform cyber adversaries, we develop new DP algorithms that simultaneously guarantee cyber resilience and privacy for the source power grids. Our algorithms

Check paper for details on:

- DP guarantees of CRO and CRO-Exp
- Connection between Bi-level and RO
- Experiment settings, data and code

Miscellaneous

Our work on energy data privacy:

1. **Synthesizing grid data with cyber resilience and privacy guarantees**
   IEEE Control Systems Letters, 2025
2. **Differentially private algorithms for synthetic power system datasets**
   IEEE Control Systems Letters, 2023
3. **Privacy-preserving convex optimization: When differential privacy meets stochastic programming**
   2025 IEEE Conference on Decision and Control
4. **Differentially private optimal power flow for distribution grids**
   IEEE Transactions on Power Systems, 2021. 🎗 Best 2019–2021 Paper Award
5. **Differentially private distributed optimal power flow**
   2020 IEEE Conference on Decision and Control

New perspective on energy data privacy via diffusion models:

Monday, October 27, 2025 at 2:57 PM – 3:09 PM EDT

Synthesizing Representative Power Flow Datasets via Constrained Diffusion Models

Part of MD31 – Grid Modeling, Simulation, and Resilience | Building B Level 2 B205

Track: Contributed

Milad Hoseinpour