Enabling Access to Power Grid Data and Models via Differential Privacy

Vladimir Dvorkin

Department of Electrical Engineering and Computer Science

University of Michigan

4th Champery Power Conference

February 12, 2026

UNIVERSITY OF MICHIGAN

Data gap between power systems and other communities

- **Open-access datasets** have fueled breakthroughs in many fields:

Computer vision
(ImageNet)

Speech recognition
(LibriSpeech)

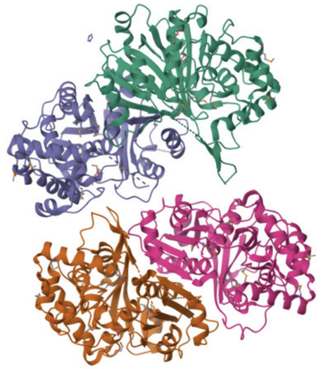

Biology
(UniProt)

- Power systems research **lags behind**:

  Real grid data are **hard to access** — due to security and regulation

- Most available datasets are **synthetic**, limiting realism and impact

- **Result:** Enthusiasm and progress in AI, optimization, and control in power systems does not convert to immediate benefits to real-world systems

Where energy data should go?

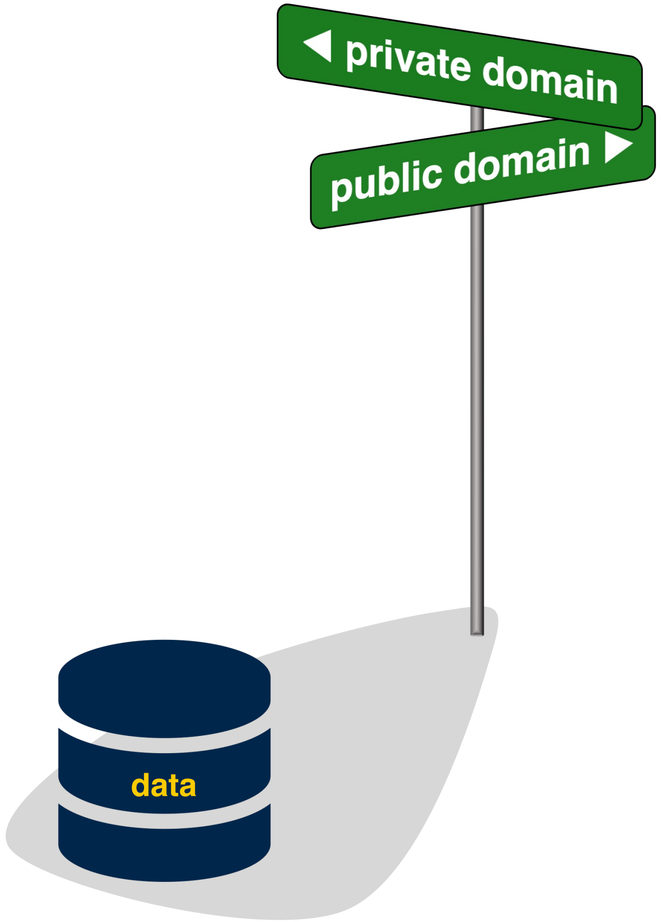

Arguments in favor of **private** data:

- Privacy and security
- Regulatory compliance
- Competitive advantage

Arguments in favor of **public** data:

- Improved decision-making
- Less barriers for entry
- Innovation, research

Is there an **non-discrete** answer to this question?

Grid data access is heavily regulated

- Typical response to data sharing request looks something like this:

  "*We will not be able to share actual grid data due to FERC CEII restrictions.*"

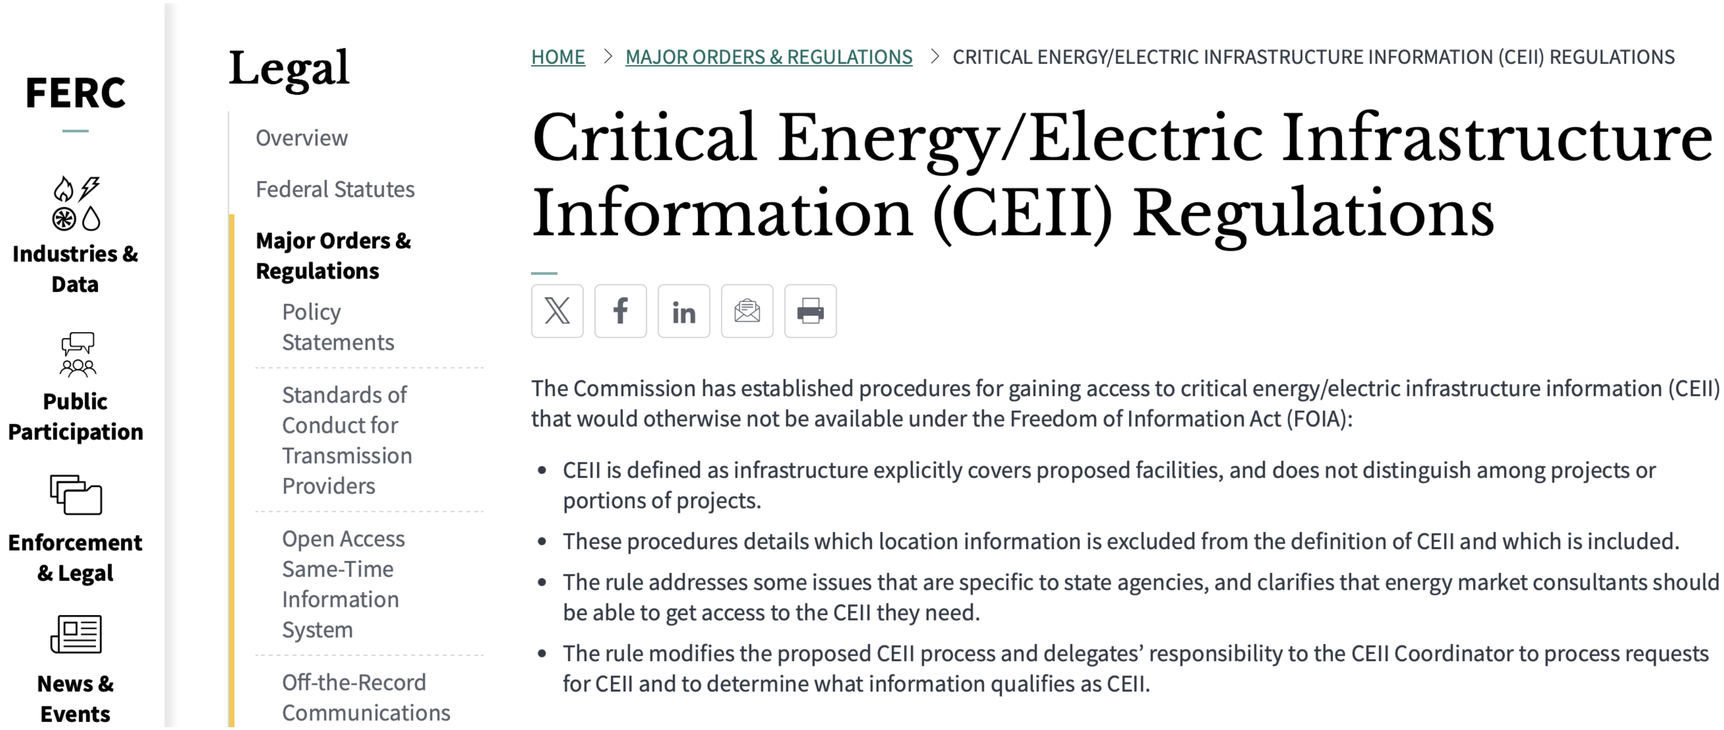

FERC

Industries & Data

Public Participation

Enforcement & Legal

News & Events

Legal

Overview

Federal Statutes

Major Orders & Regulations

Policy Statements

Standards of Conduct for Transmission Providers

Open Access Same-Time Information System

Off-the-Record Communications

HOME > MAJOR ORDERS & REGULATIONS > CRITICAL ENERGY/ELECTRIC INFRASTRUCTURE INFORMATION (CEII) REGULATIONS

Critical Energy/Electric Infrastructure Information (CEII) Regulations

The Commission has established procedures for gaining access to critical energy/electric infrastructure information (CEII) that would otherwise not be available under the Freedom of Information Act (FOIA):

- CEII is defined as infrastructure explicitly covers proposed facilities, and does not distinguish among projects or portions of projects.
- These procedures details which location information is excluded from the definition of CEII and which is included.
- The rule addresses some issues that are specific to state agencies, and clarifies that energy market consultants should be able to get access to the CEII they need.
- The rule modifies the proposed CEII process and delegates' responsibility to the CEII Coordinator to process requests for CEII and to determine what information qualifies as CEII.

- Grid engineers are not trained in data transparency and management to think how to overcome such barriers
- Vendors of power systems components are hesitant to disclose proprietary parameters and models

Synthetic data is a viable alternative to real data sharing

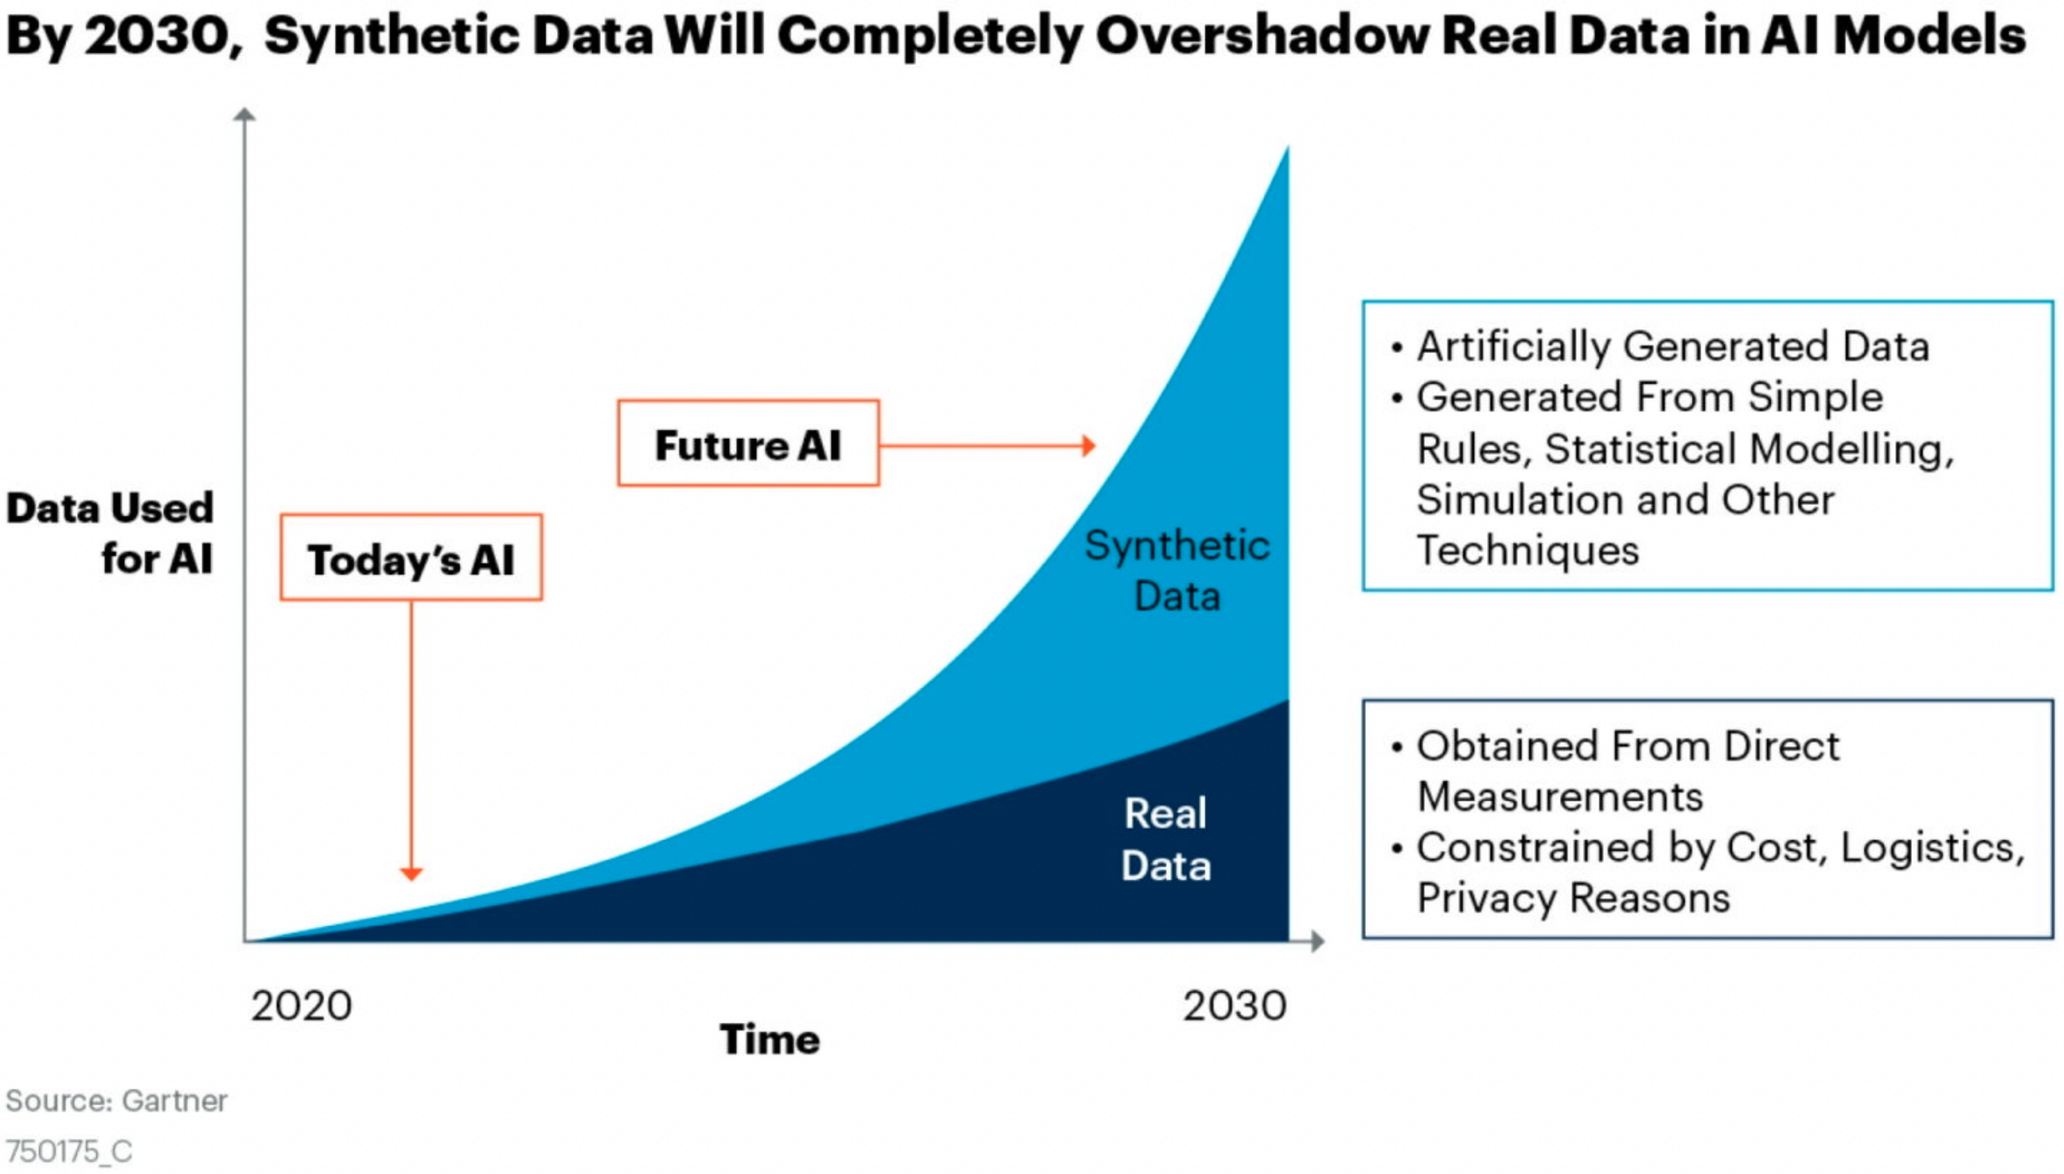

Synthetic power systems datasets

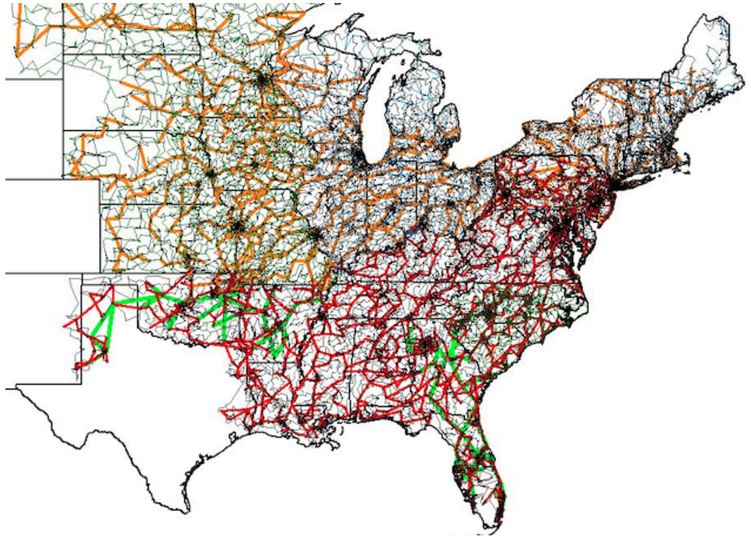

Texas A&M University Grid Datasets (from 37 to 80k+ bus networks)

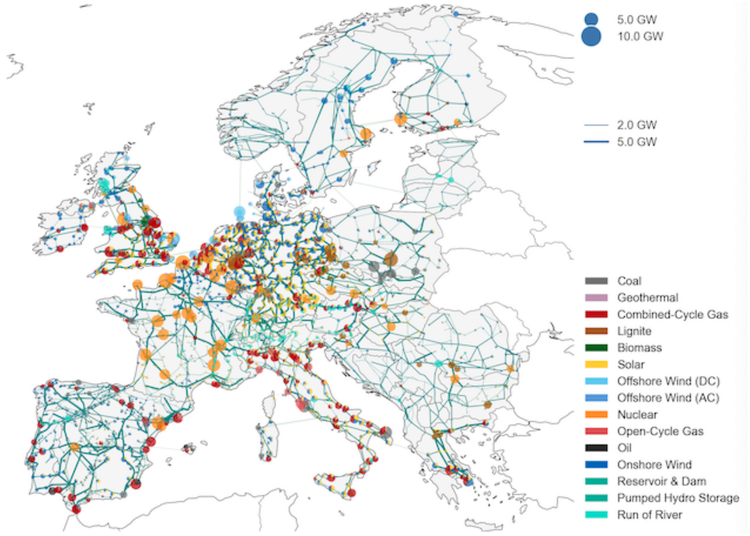

PyPSA-Eur: synthetic dataset of Europe covering the full ENTSO-E area

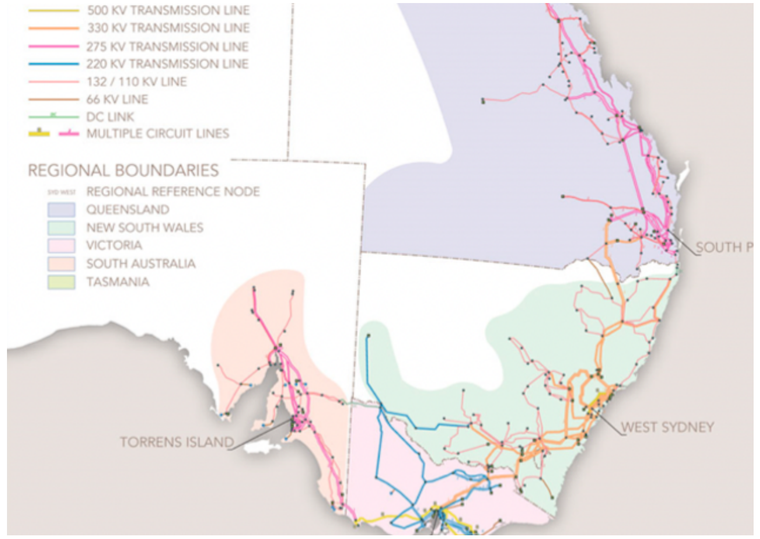

Synthetic Data of the National Electricity Market (Australia)

Why these datasets may not satisfy our needs?

- "[...] data bears **no relation** to the actual grid [...] except that generation and load profiles are similar, based on public data"
- "This test case represents a synthetic (**fictitious**) transmission"
- "This case is synthetic and **does not** model the actual grid"

Our work

We develop algorithms to synthesize credible synthetic datasets from real power systems, while controlling privacy and security risks

- **Algorithmic formalization of trust** in energy data sharing
  - Moving from subjective decisions to quantitative guarantees
- **Private or public grid data?** Our algorithms provide a *non-discrete* solution to this dilemma
  - Spectrum of privacy-utility tradeoffs, not just share/don't share
- **Comprehensive data scope:**
  - Optimization datasets (OPF, unit commitment, economic dispatch)
  - Dynamic models (e.g., grid frequency dynamics)
  - Operational records (e.g., power flow datasets)

Our algorithms enable controlled grid data disclosures while resolving privacy, security, and logistical concerns

Agenda

1. Intro
2. Formalization of energy data privacy
3. Synthesizing optimization data with privacy and cyber security guarantees
4. Synthesizing power system dynamics models with privacy guarantees
5. Outro

Formalizing differential privacy (DP)

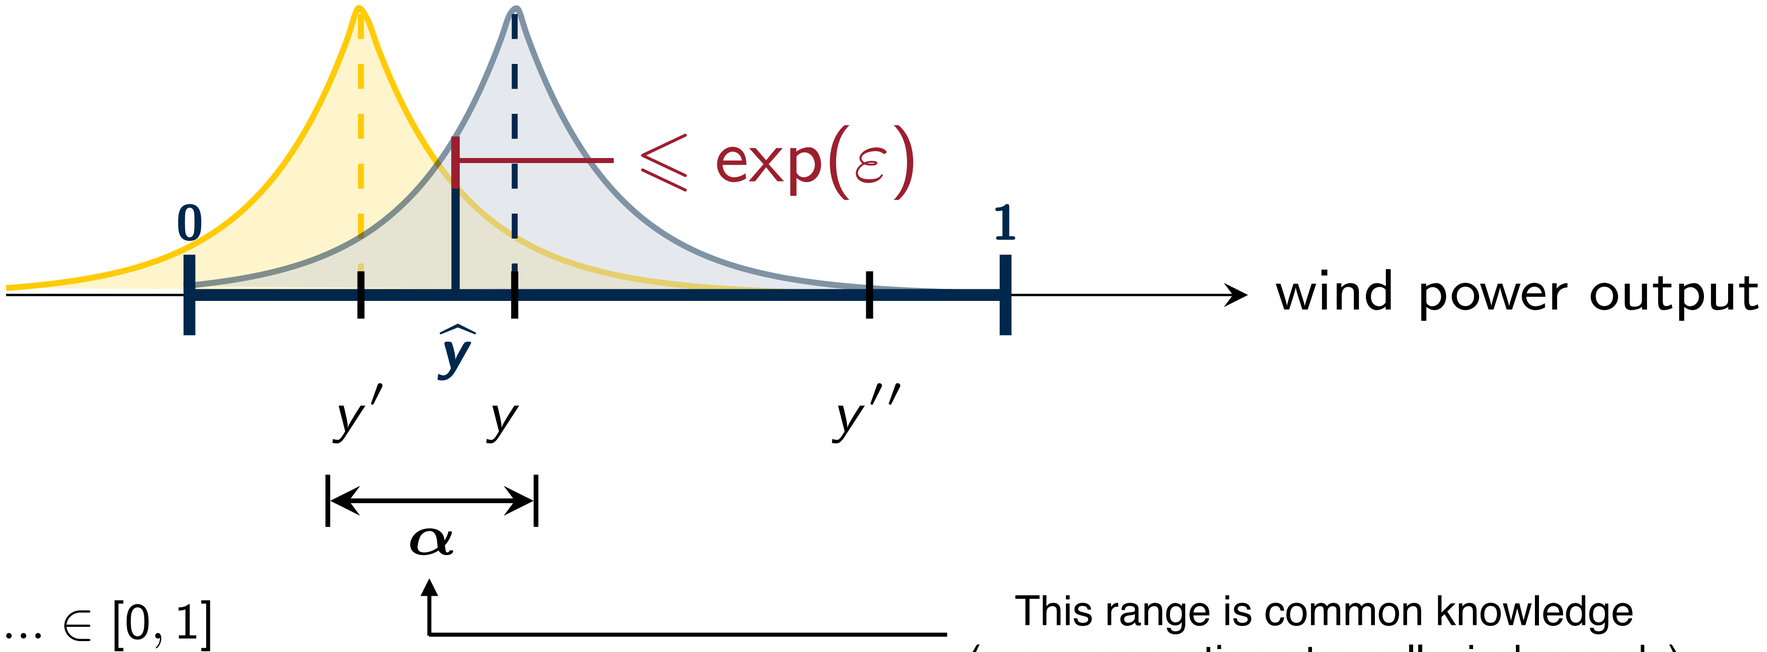

- Wind power records $y, y', y'', \ldots \in [0, 1]$
- For given $\alpha > 0$, records y and y' are $\alpha-$adjacent if $\|y - y'\| \leqslant \alpha$

This range is common knowledge
(e.g., generation at small wind speeds)

- Let Lap(α/ε) be a zero-mean random Laplacian noise
- The synthetic counterpart of real record y is $\widehat{\boldsymbol{y}}(y) = y + \text{Lap}(\alpha/\varepsilon)$
- For some parameter $\varepsilon > 0$, the release $\widehat{\boldsymbol{y}}(y)$ is $\varepsilon-$DP if

$$\text{for any } \alpha\text{−adjacent pair } (y, y'): \quad \left| \log \left(\frac{\mathbb{P}[\,\widehat{\boldsymbol{y}}(y\,)\,]}{\mathbb{P}[\,\widehat{\boldsymbol{y}}(y')\,]} \right) \right| \leqslant \varepsilon \qquad \text{(bounded log-likelihood ratio)}$$

Key principle of data privacy: perturb data but preserve its value

The Starry Night by Vincent Van Gogh

1888 (Musée d'Orsay's, Paris)

1889 (Museum of Modern Art, NYC)

The value of each painting is well over $100 million

Key principle of data privacy: perturb data but preserve its value

The Starry Night by Vincent Van Gogh

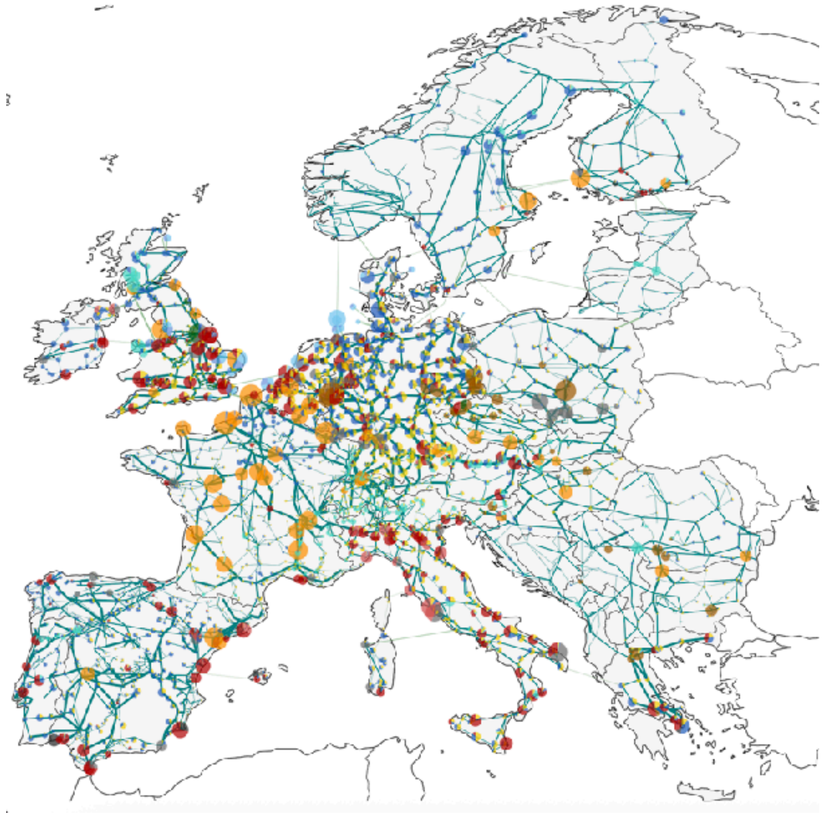

1888 (Musée d'Orsay's, Paris)

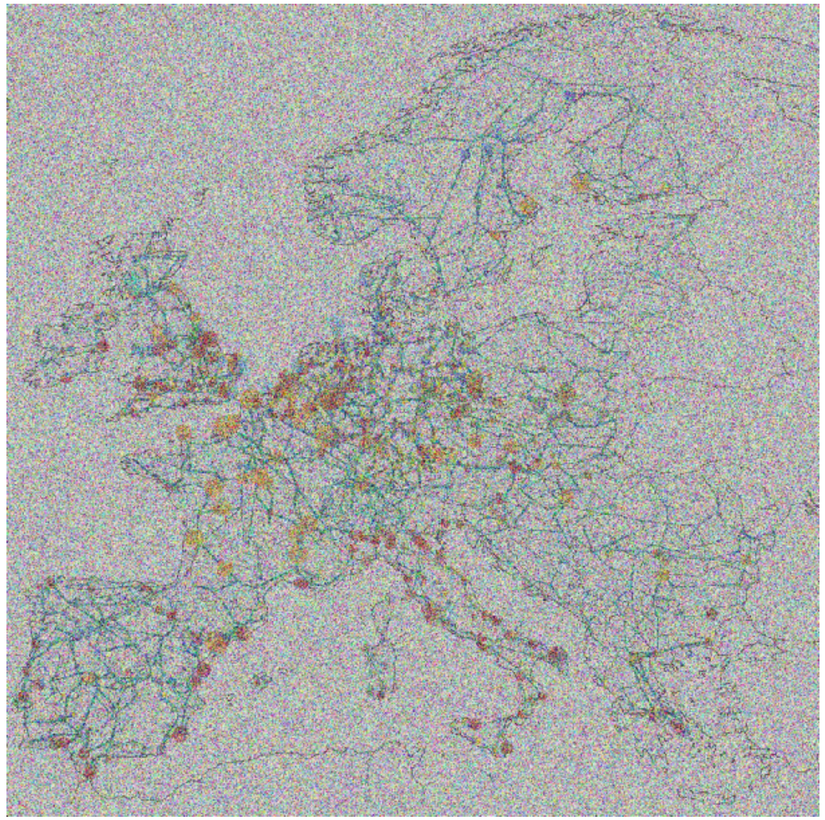

1889 (Museum of Modern Art, NYC)

The value of each painting is well over $100 million

Differential privacy (DP) enables controlled disclosure of grid data

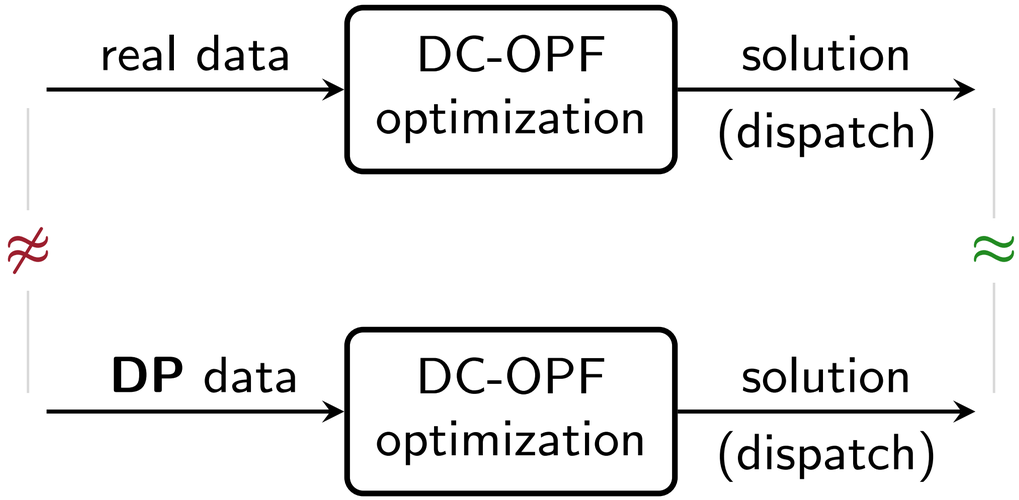

- DP principle: obfuscate real data (add noise) but preserve its value
- In the DC-OPF setting: obfuscate grid data but preserve the OPF solution
- Formal *privacy guarantee*: released DP data does NOT disclose the real data
- Many applications to synthesizing high-quality transmission, load and generation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

Example: Synthesizing DP transmission data for optimal power flow 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

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

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

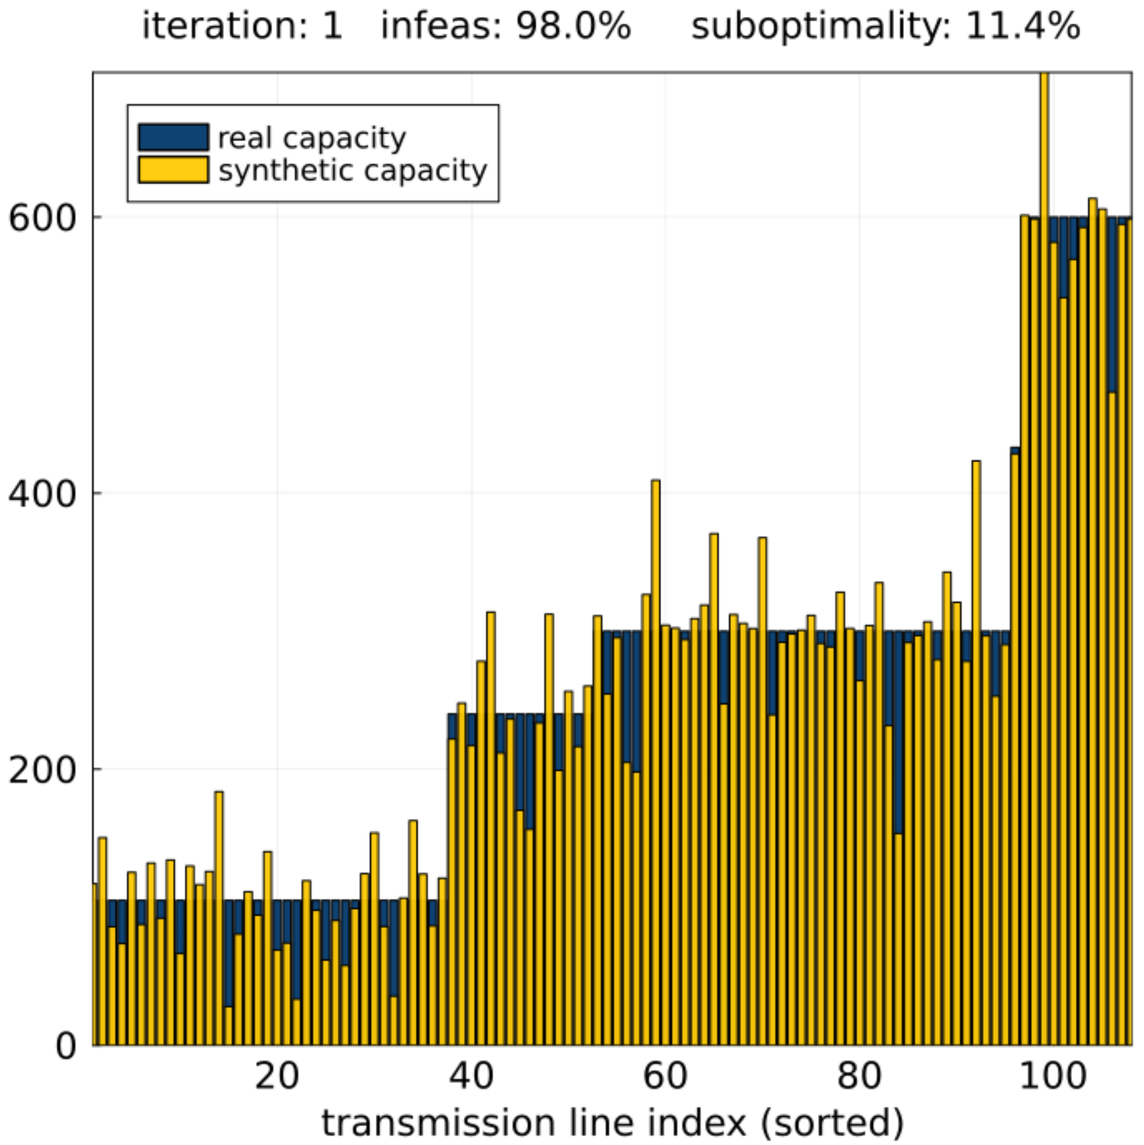

Dvorkin, V., Botterud A. Differentially private algorithms for synthetic power system datasets, IEEE Control Systems Letters, 2023.

Agenda

1. Intro
2. Formalization of energy data privacy
3. **Synthesizing optimization data with privacy and cyber security guarantees**
4. Synthesizing power system dynamics models with privacy guarantees
5. Outro

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:

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

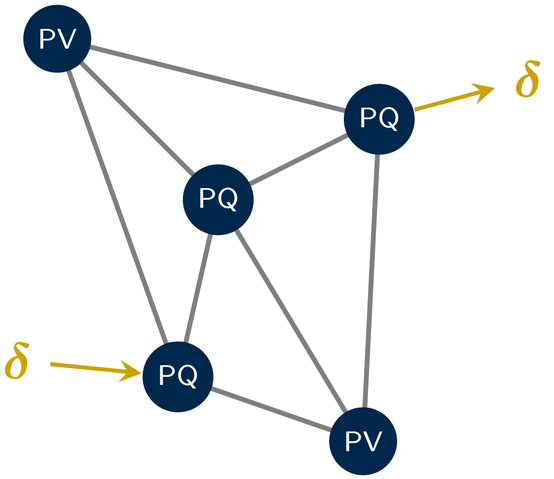

- The problem seeks a *stealthy* attack vector $\boldsymbol{\delta}$ that maximizes the OPF cost

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

Numerical evidence

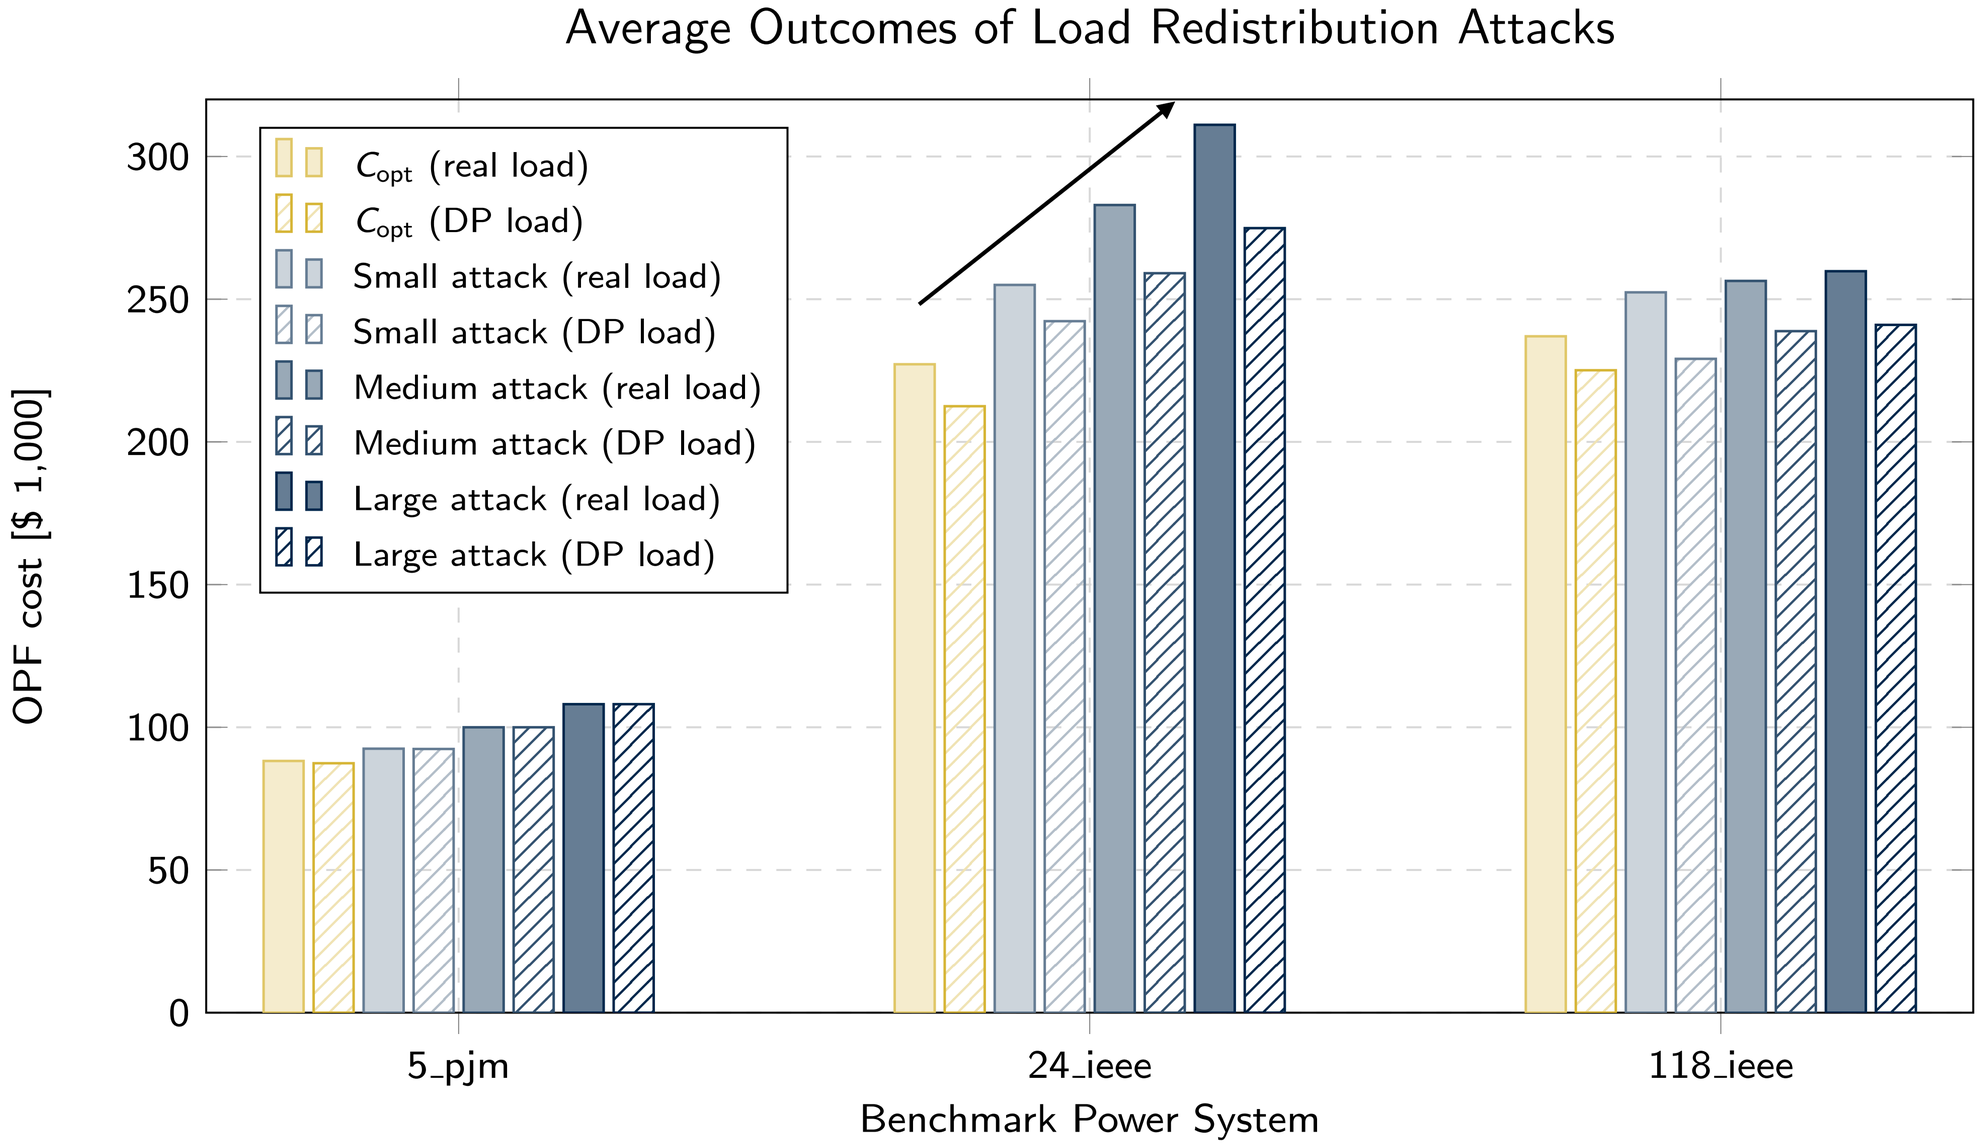

Post-processing optimization to minimize attack damage

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

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

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

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 $\mathcal{K}'$

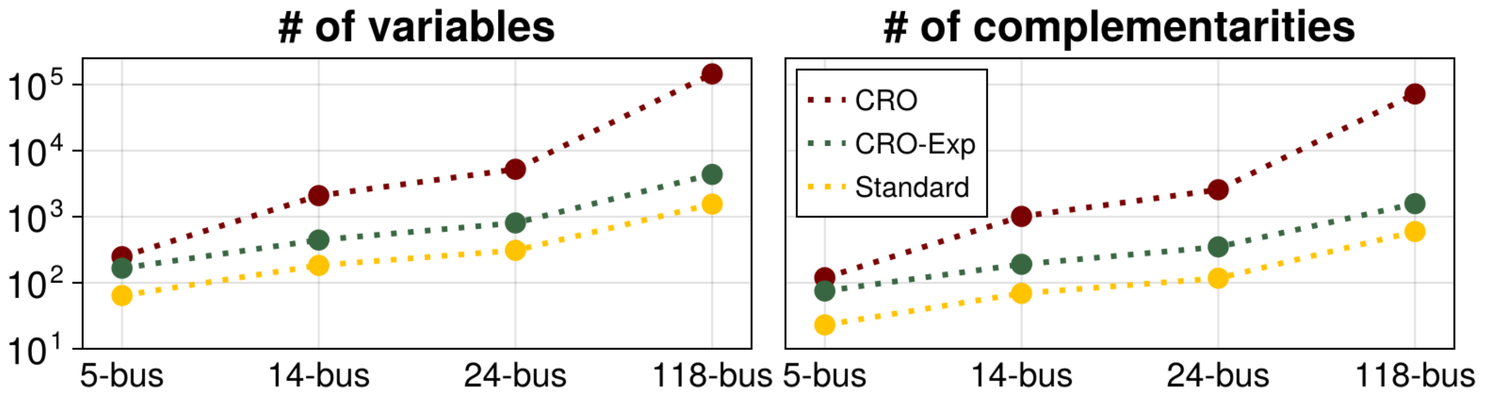

Selecting only important constraints substantially reduces the computational burden

CRO application to PJM 5-Bus system

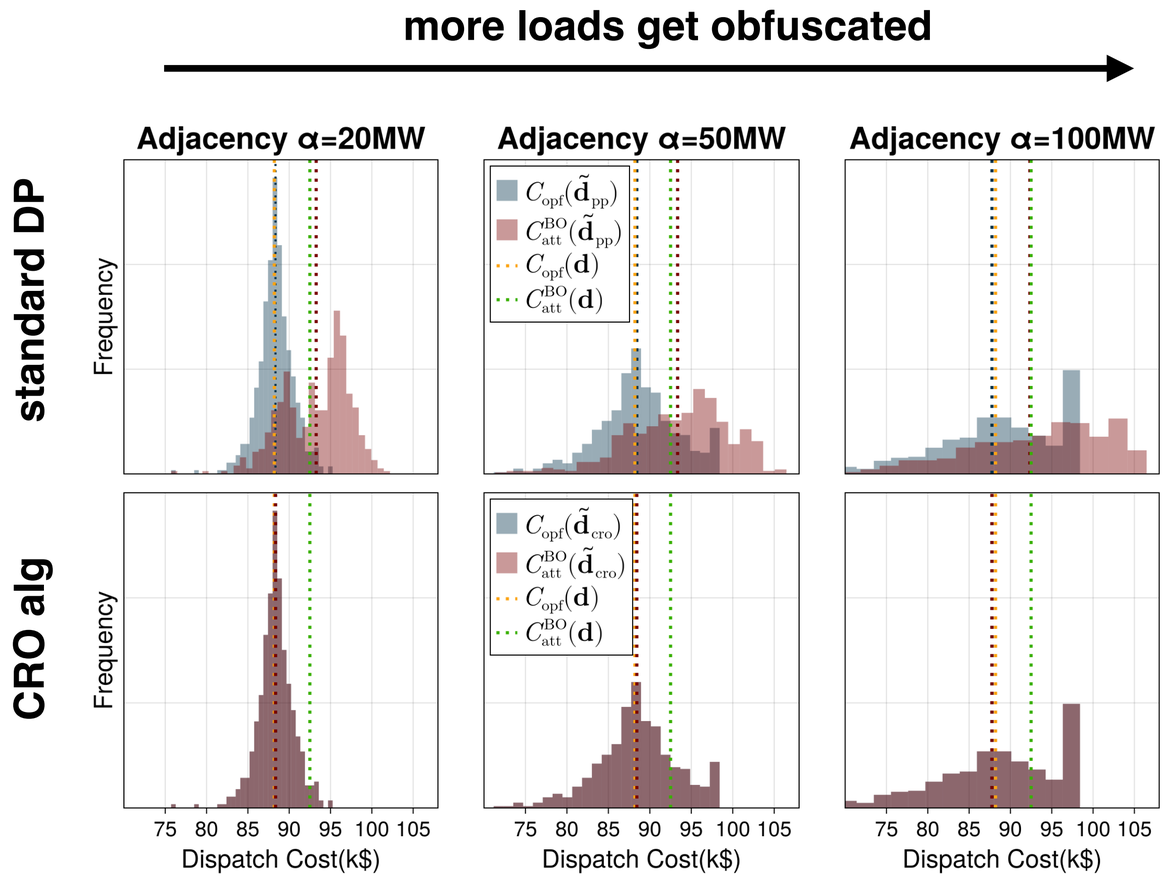

OPF cost distributions in normal and post-attack operation

- 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

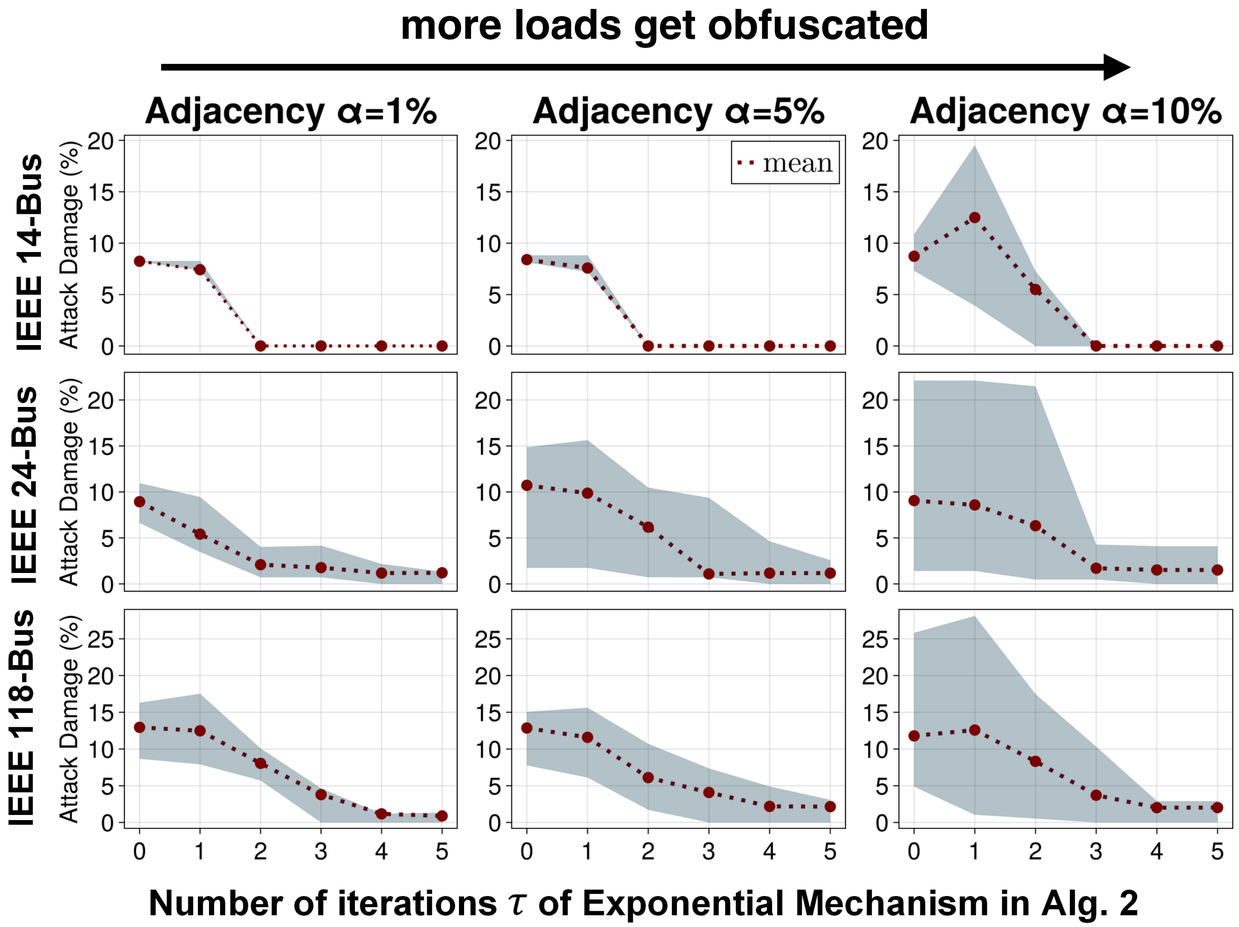

Post-attack damage as a function of constraints selected for post-processing optimization

- 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**.

Preliminary 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 bi-level and robust optimization techniques)

438 IEEE CONTROL SYSTEMS LETTERS, VOL. 9, 2025

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

Agenda

1. Intro
2. Formalization of energy data privacy
3. Synthesizing optimization data with privacy and cyber security guarantees
4. **Synthesizing power system dynamics models with privacy guarantees**
5. Outro

The Challenge

- **Stability** of power grids is critical

  Renewables, inverter-based resources, and data centers reshape grid dynamics

- We need a **shared dynamical model** so independent parties can design optimal control

- But releasing the real model **exposes sensitive system details**

- **Goal:** Synthesize a model that behaves nearly identically — without disclosing the original model

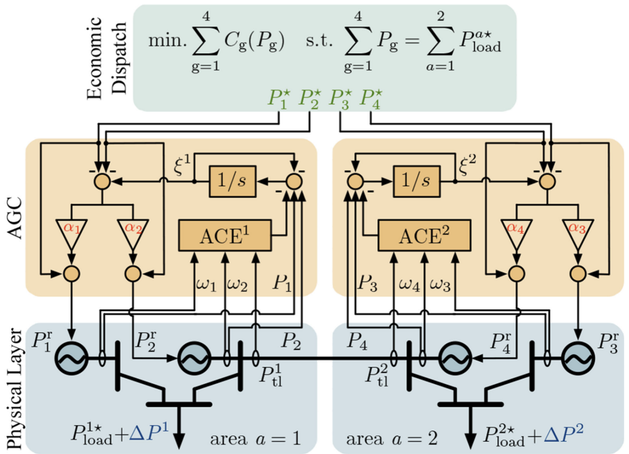

picture courtesy of Sairaj Dhople

Grid frequency dynamic model

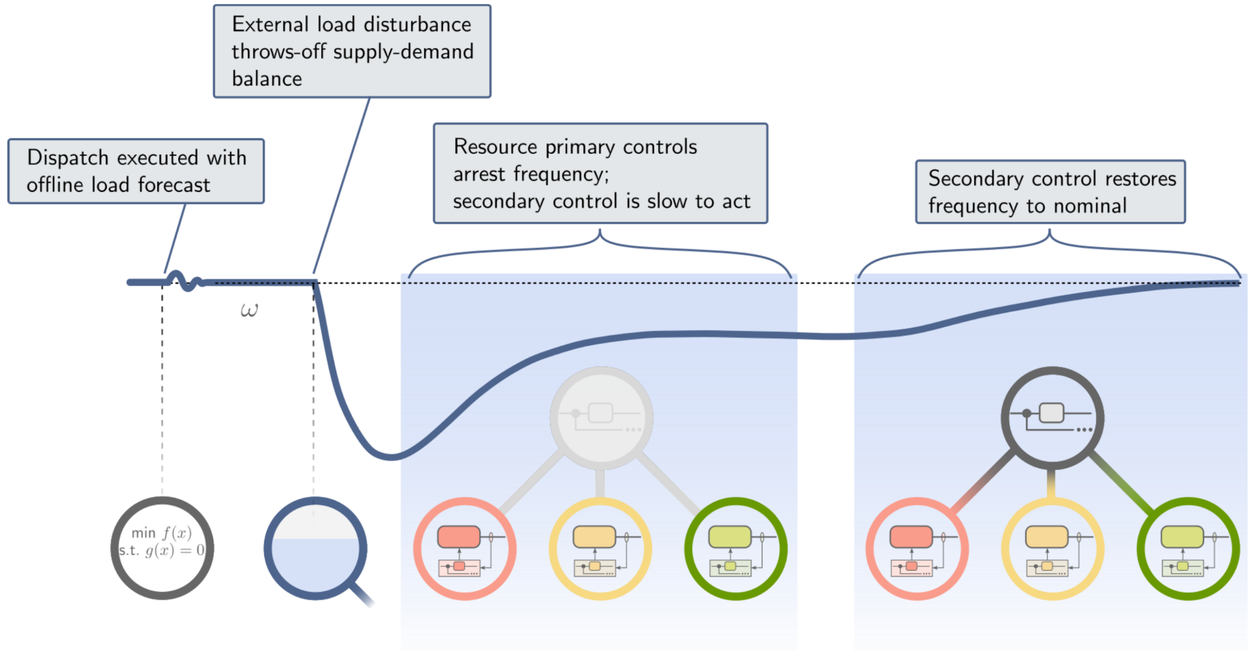

picture courtesy of Sairaj Dhople

topology/admittance

$$\dot{\omega} = \mathbf{M}^{-1}(\mathbf{K}\delta + \mathbf{p} - \mathbf{d} - \mathbf{D}\omega)$$

inertia

damping

swing equation

$$\dot{\delta} = \omega$$

phase angle dynamics

$$\dot{\mathbf{p}} = -\frac{1}{\mathbf{T}}(\mathbf{p} - \mathbf{r} + \mathbf{R}\omega)$$

generator control

How to release the parameters of system dynamics in a privacy-preserving way?

Classic solution based on Kron reduction

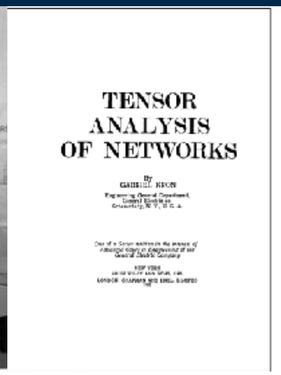
TENSOR
ANALYSIS
OF NETWORKS

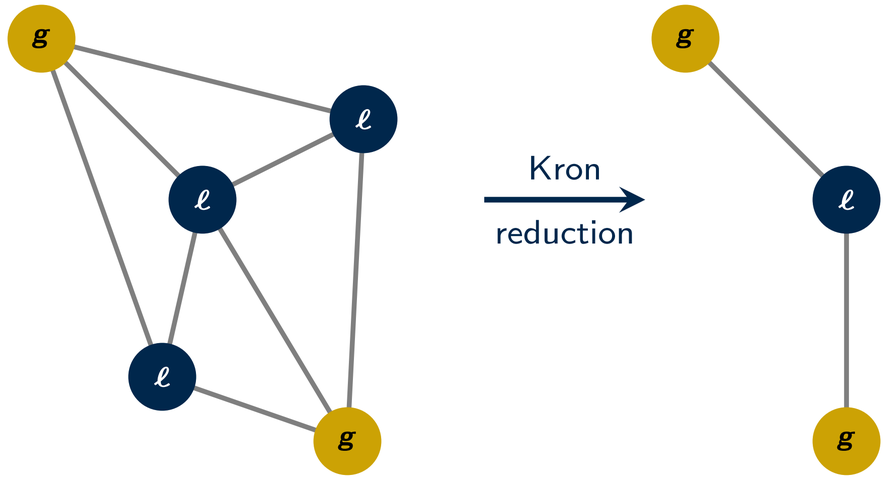

- Partitioning of the swing equation:

$$\begin{bmatrix} \mathbf{M}_g & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{\omega}_g \\ \boldsymbol{\omega}_\ell \end{bmatrix} = \begin{bmatrix} \mathbf{K}_{gg} & \mathbf{K}_{g\ell} \\ \mathbf{K}_{\ell g} & \mathbf{K}_{\ell\ell} \end{bmatrix} \begin{bmatrix} \boldsymbol{\delta}_g \\ \boldsymbol{\delta}_\ell \end{bmatrix} + \begin{bmatrix} \mathbf{p}_g \\ \mathbf{0} \end{bmatrix} - \begin{bmatrix} \mathbf{d}_g \\ \mathbf{d}_\ell \end{bmatrix} - \begin{bmatrix} \mathbf{D}_g & \mathbf{0} \\ \mathbf{0} & \mathbf{D}_\ell \end{bmatrix} \begin{bmatrix} \boldsymbol{\omega}_g \\ \boldsymbol{\omega}_\ell \end{bmatrix}$$

- Kron-reduced swing equation with the identical dynamics

$$\mathbf{M}_g \dot{\boldsymbol{\omega}}_g = \mathbf{K}_{\text{red}} \boldsymbol{\delta}_g - \mathbf{D}_g \boldsymbol{\omega}_g + (\mathbf{p} - \mathbf{d}_g) - \mathbf{K}_{\text{ac}} \mathbf{d}_\ell$$

- Does the release of the reduced equation preserves privacy?
  No! [SH Low (2024); Deka, Kekatos & Cavraro (2024)]

Differentially private obfuscation of the frequency dynamic model

- **Step 1:** Perturb dynamic model with random Laplace noise:

$$\tilde{\mathbf{M}} = \mathbf{M} + \text{Lap}(\alpha_{\mathbf{M}}/\varepsilon_{\mathbf{M}}), \quad \tilde{\mathbf{K}} = \mathbf{K} + \text{Lap}(\alpha_{\mathbf{K}}/\varepsilon_{\mathbf{K}}), \quad \tilde{\mathbf{D}} = \mathbf{D} + \text{Lap}(\alpha_{\mathbf{D}}/\varepsilon_{\mathbf{D}})$$

- **Step 2:** Post-process the perturb parameters

$$\underset{\mathbf{M},\mathbf{K},\mathbf{D} \cup \boldsymbol{\omega},\boldsymbol{\delta},\mathbf{p}}{\text{minimize}} \quad \int_0^T \underbrace{\left\|\boldsymbol{\omega}(t|\mathbf{M},\mathbf{K},\mathbf{D}) - \boldsymbol{\omega}_{\text{pub}}(t)\right\|_2^2}_{\text{model fidelity}} dt + \underbrace{\lambda\left\|\mathbf{M} - \tilde{\mathbf{M}}\right\|_2^2}_{\text{regularization}} + \dots$$

$$\text{subject to} \quad \mathbf{M}\dot{\boldsymbol{\omega}} = \mathbf{K}_{\text{red}}\boldsymbol{\delta} - \mathbf{D}\boldsymbol{\omega} + (\mathbf{p} - \mathbf{d}_g) - \mathbf{K}_{\text{ac}}\mathbf{d}_\ell$$

$$\dots$$

- The optimal solution $\mathbf{M}, \mathbf{K}, \mathbf{D}$ is optimized to recover the public frequency dynamic
- Solved over iterations using the combination of the gradient descent and adjoint method
- **Privacy guarantee:** $\mathbf{M}, \mathbf{K}, \mathbf{D}$ preserves $\varepsilon-$DP for $\alpha-$adjacent model parameters when setting

$$\varepsilon_{\mathbf{M}} = \varepsilon_{\mathbf{K}} = \varepsilon_{\mathbf{D}} = \frac{1}{3}\varepsilon$$

Example: IEEE 30-bus system

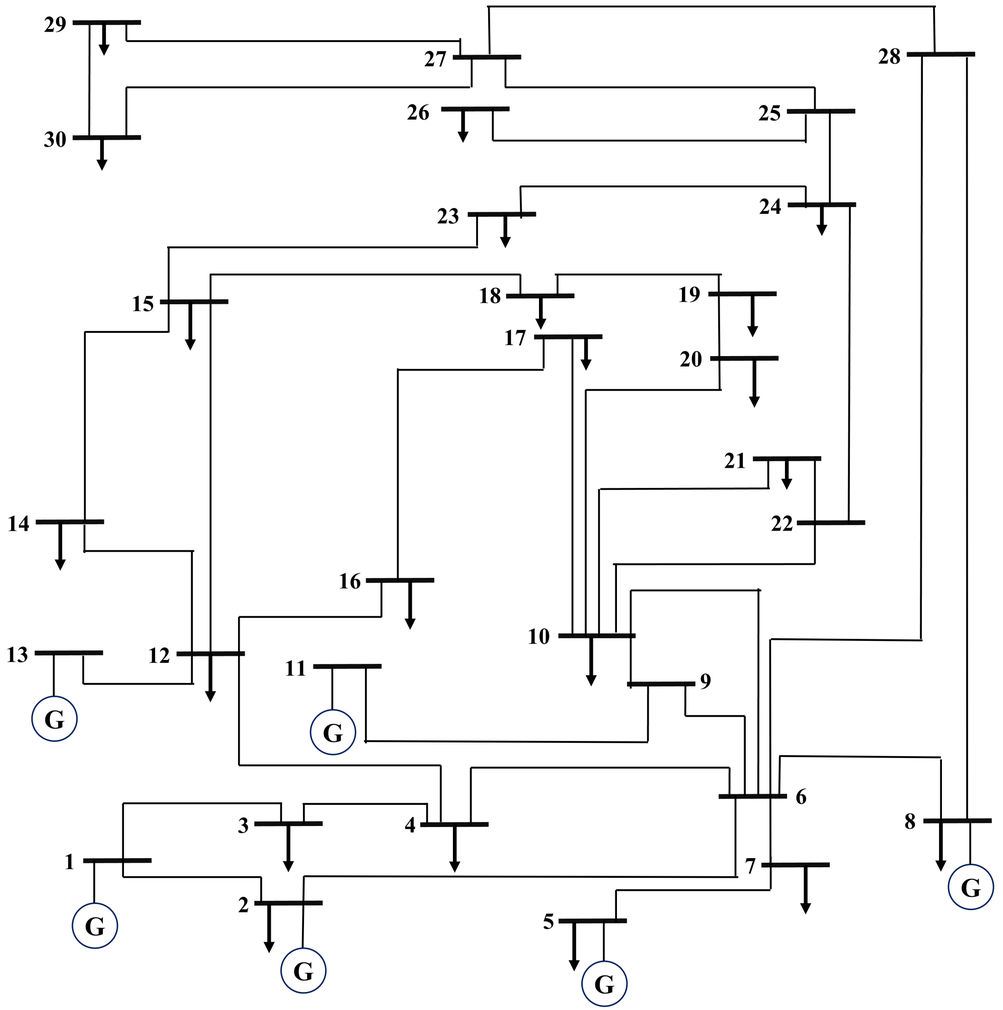

Kron reduction $\Longrightarrow$

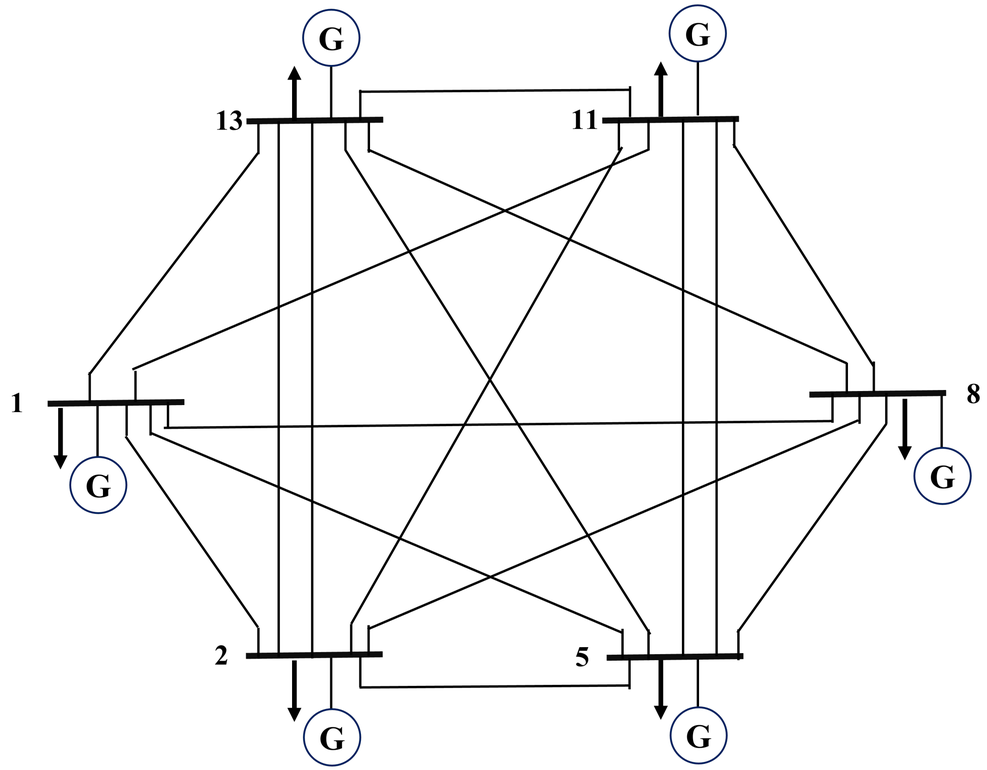

- Kron reduction from 30 to 6 buses
- DP parameters $\alpha = 1\text{pu}, \quad \varepsilon = \{2, 1, 0.5\}$

- Synthetic inertia $\quad \tilde{\mathbf{M}} = \mathbf{M} + \text{Lap}(1/\varepsilon)$
- Synthetic topology $\tilde{\mathbf{K}} = \mathbf{K} + \text{Lap}(1/\varepsilon)$
- Synthetic damping $\tilde{\mathbf{D}} = \mathbf{D} + \text{Lap}(1/\varepsilon)$

Loss function over iterations

$$L_{\text{pp}} = \frac{1}{10} \sum_{s=1}^{10} \int_0^{30} \|\boldsymbol{\omega}_{\text{pp}}^s(t) - \boldsymbol{\omega}_{\text{ref}}^s(t)\|_2^2 \, \mathrm{d}t$$

$$L_{\text{lap}} = \frac{1}{10} \sum_{s=1}^{10} \int_0^{30} \|\boldsymbol{\omega}_{\text{lap}}^s(t) - \boldsymbol{\omega}_{\text{ref}}^s(t)\|_2^2 \, \mathrm{d}t$$

The fidelity gap

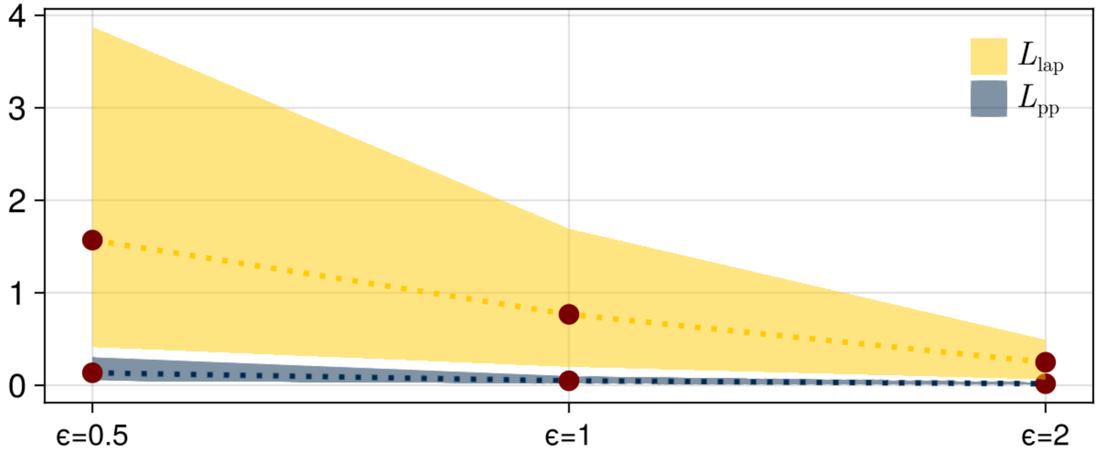

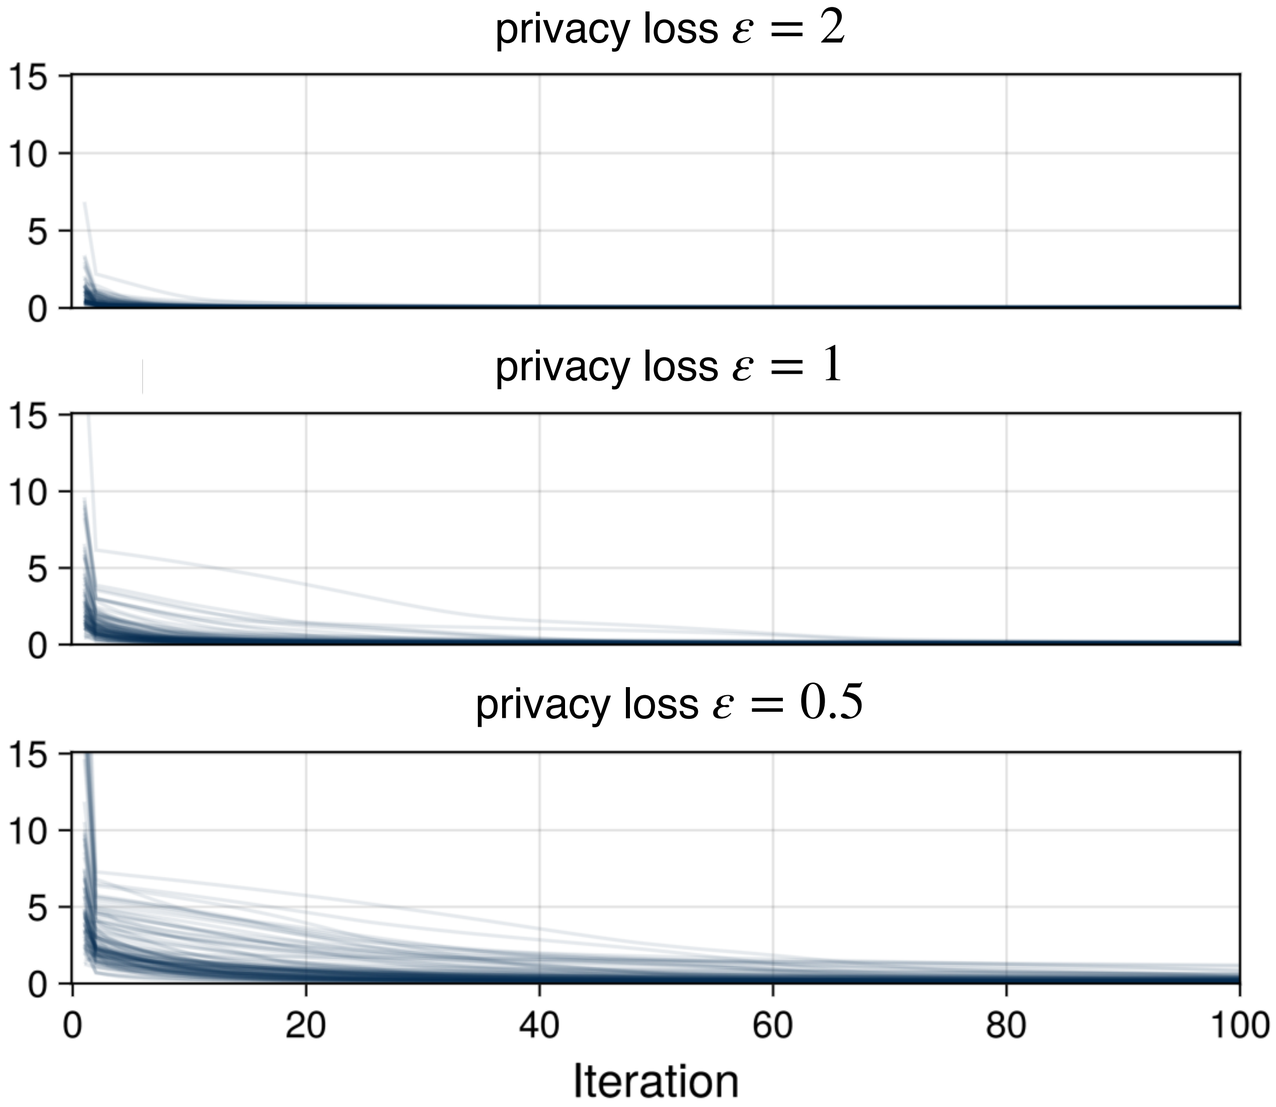

Improved fidelity with the same amount of privacy-preserving noise

Laplace mechanism vs. post-processing: Distribution of synthetic parameters

$$\mathbf{M}\dot{\boldsymbol{\omega}} = \mathbf{K}_{\text{red}}\boldsymbol{\delta} - \mathbf{D}\boldsymbol{\omega} + (\mathbf{p} - \mathbf{d}_g) - \mathbf{K}_{\text{ac}}\mathbf{d}_\ell$$

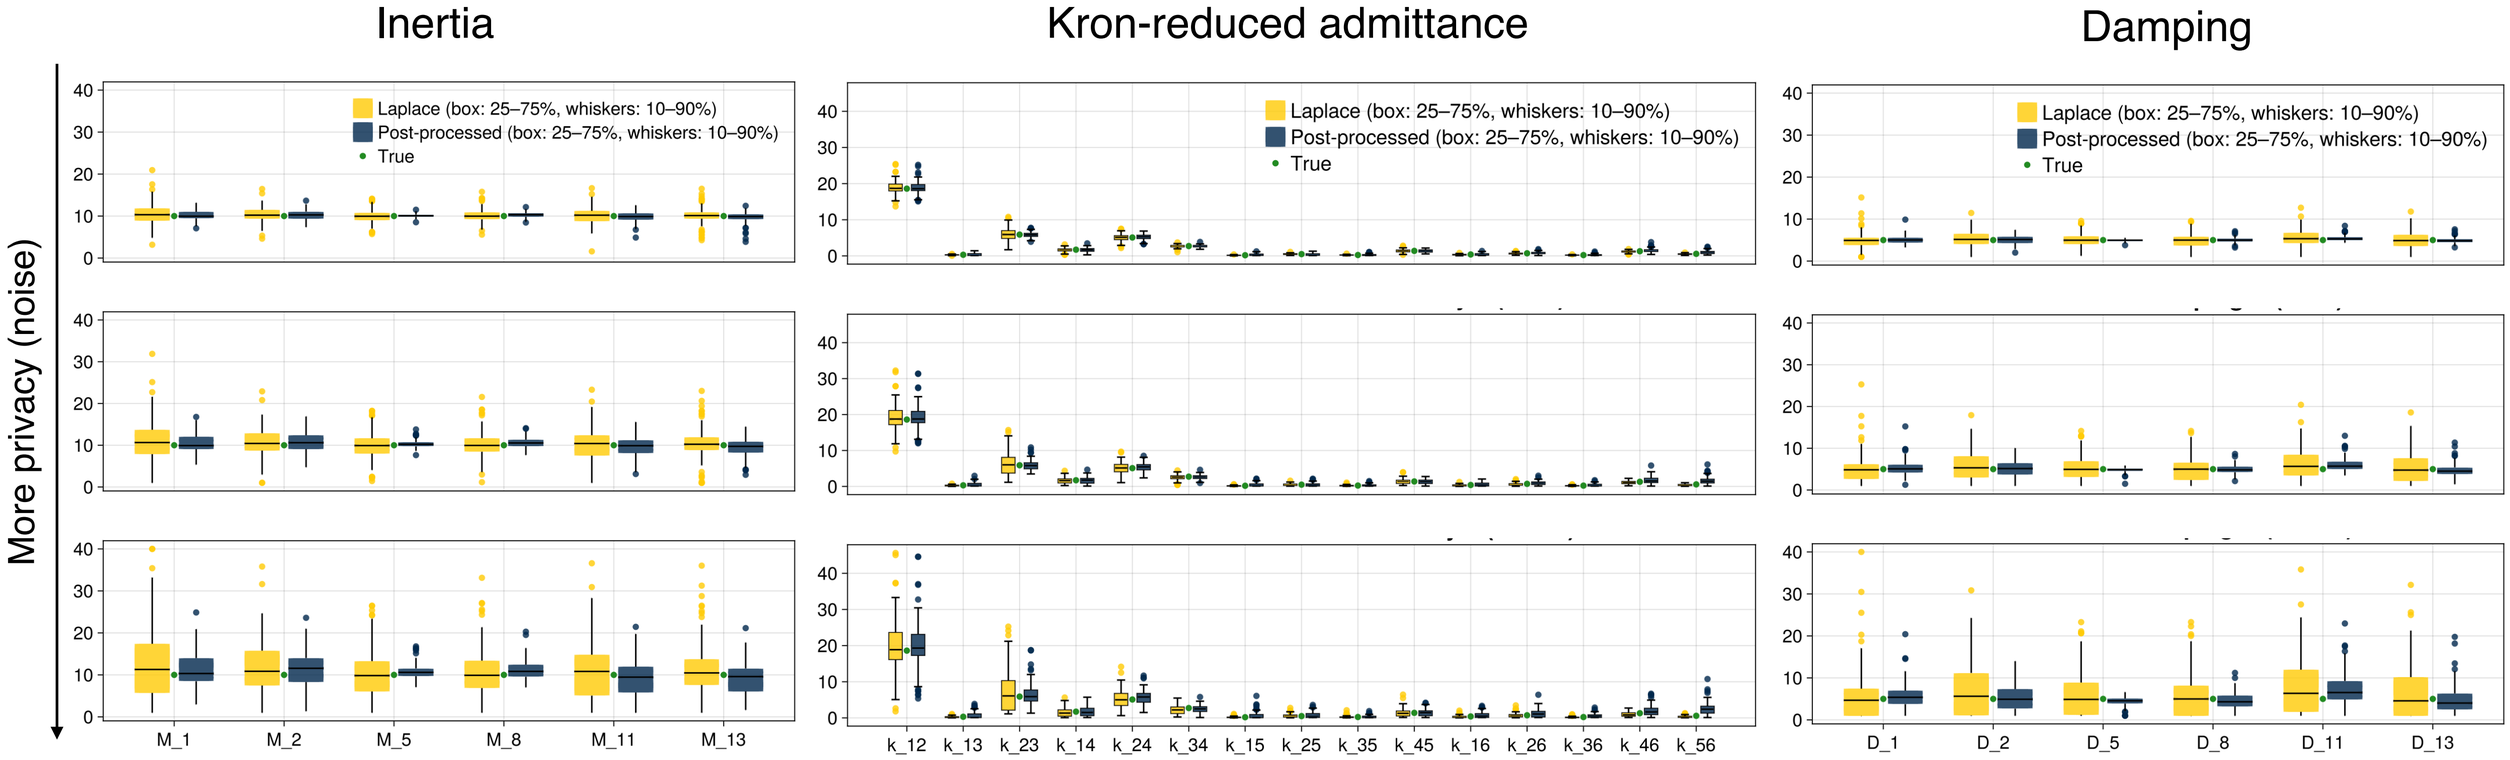

Laplace mechanism vs. post-processing: Distribution of synthetic parameters

$$\mathbf{M}\dot{\boldsymbol{\omega}} = \mathbf{K}_{\text{red}}\boldsymbol{\delta} - \mathbf{D}\boldsymbol{\omega} + (\mathbf{p} - \mathbf{d}_g) - \mathbf{K}_{\text{ac}}\mathbf{d}_\ell$$

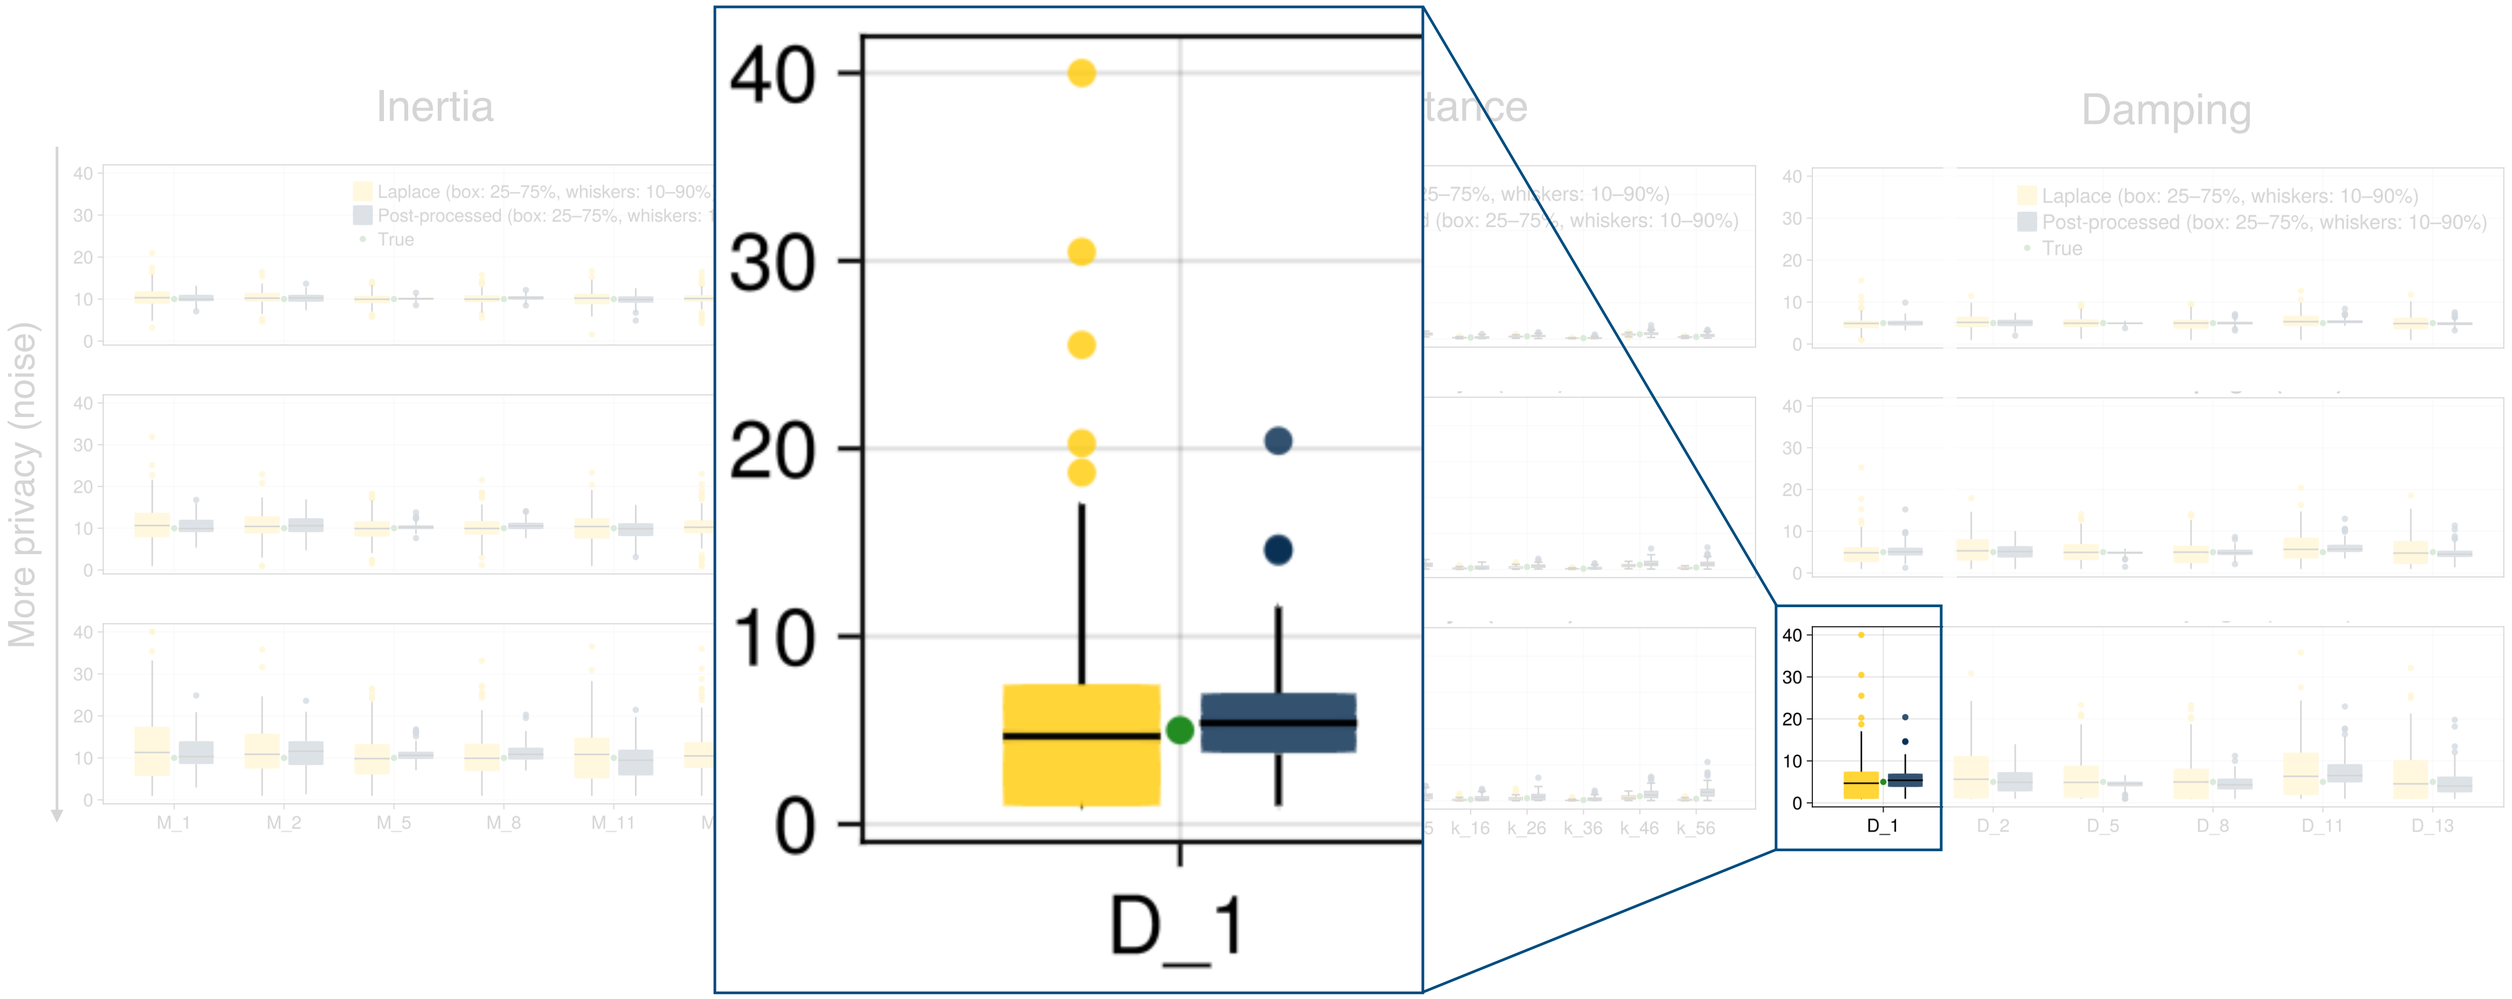

Laplace mechanism vs. post-processing: Model fidelity and resulting dynamics

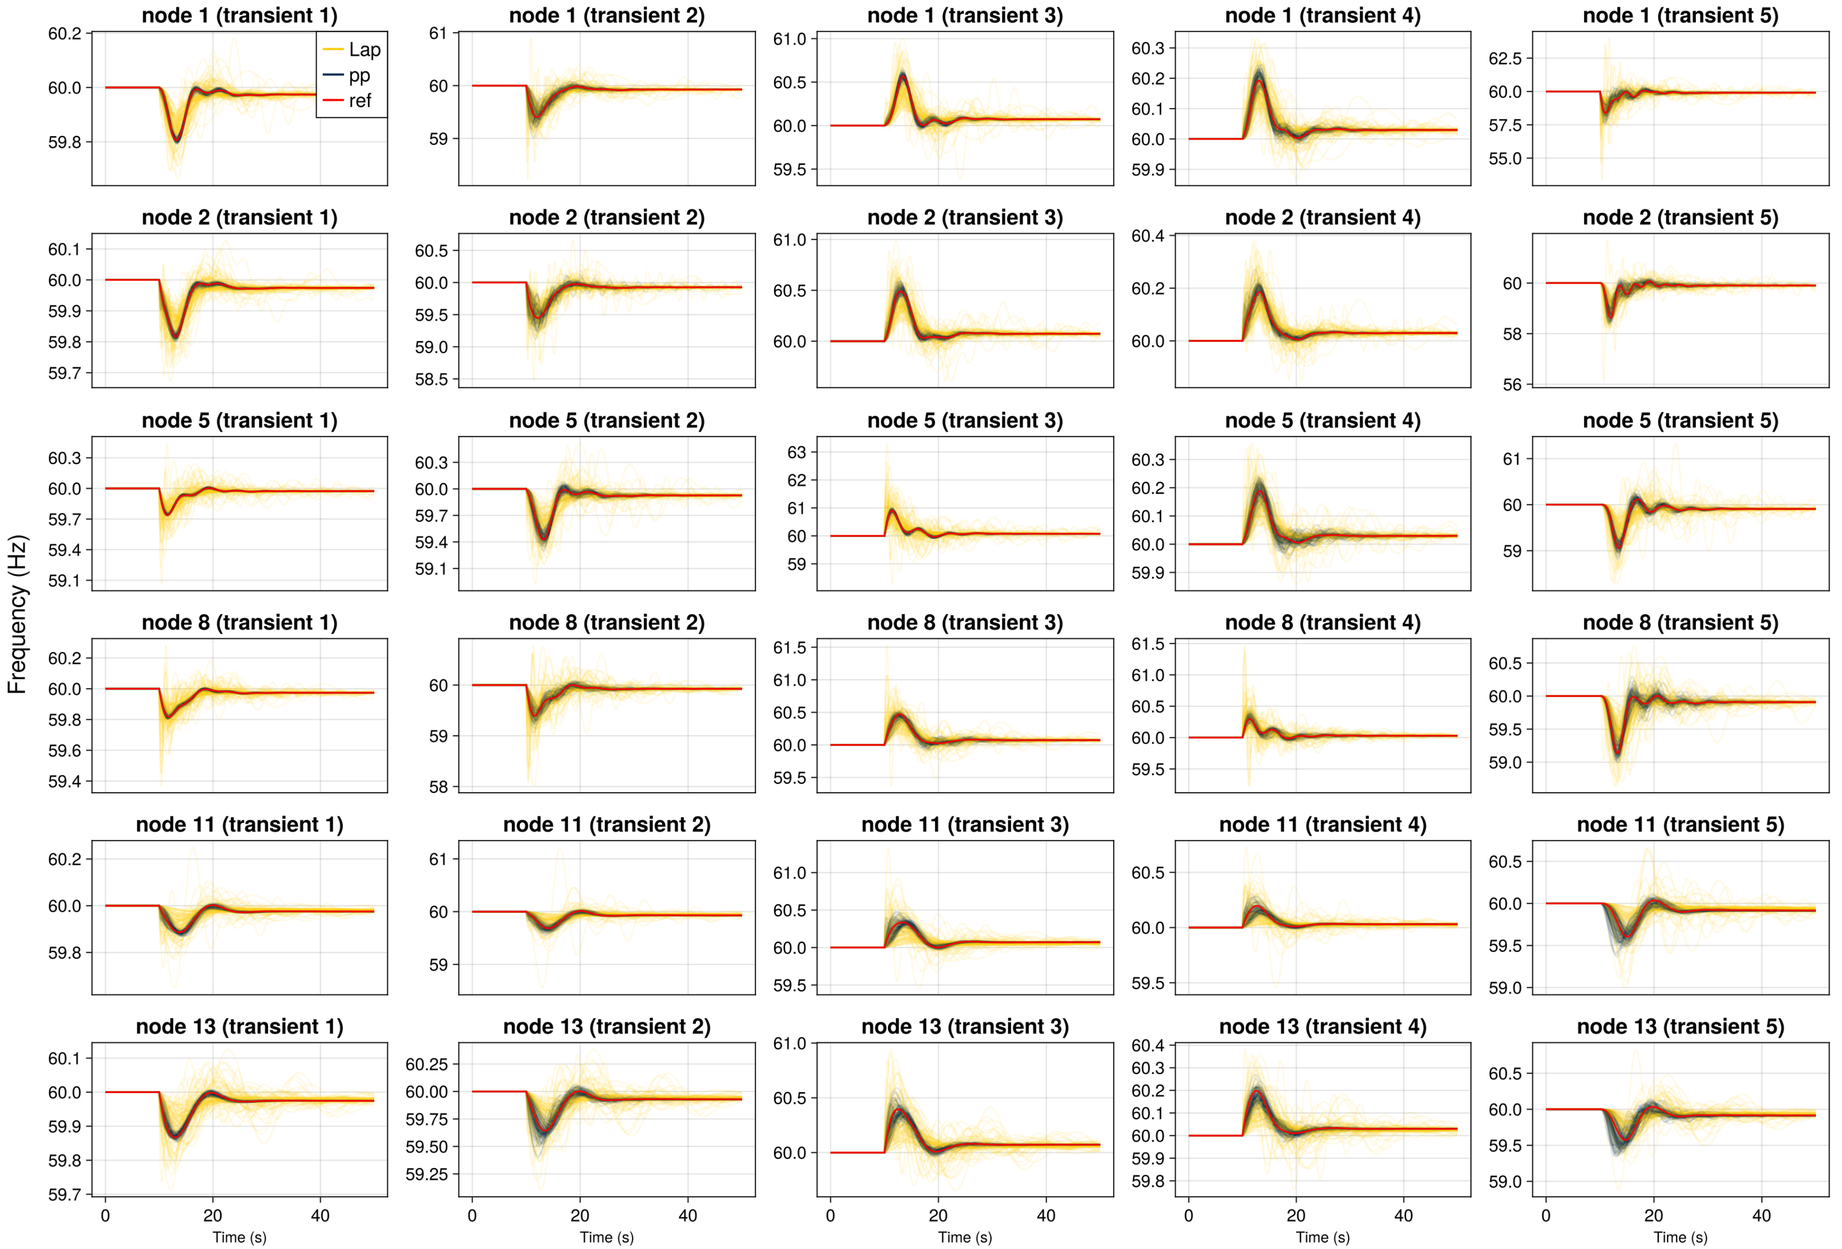

Laplace mechanism vs. post-processing: Model fidelity and resulting dynamics

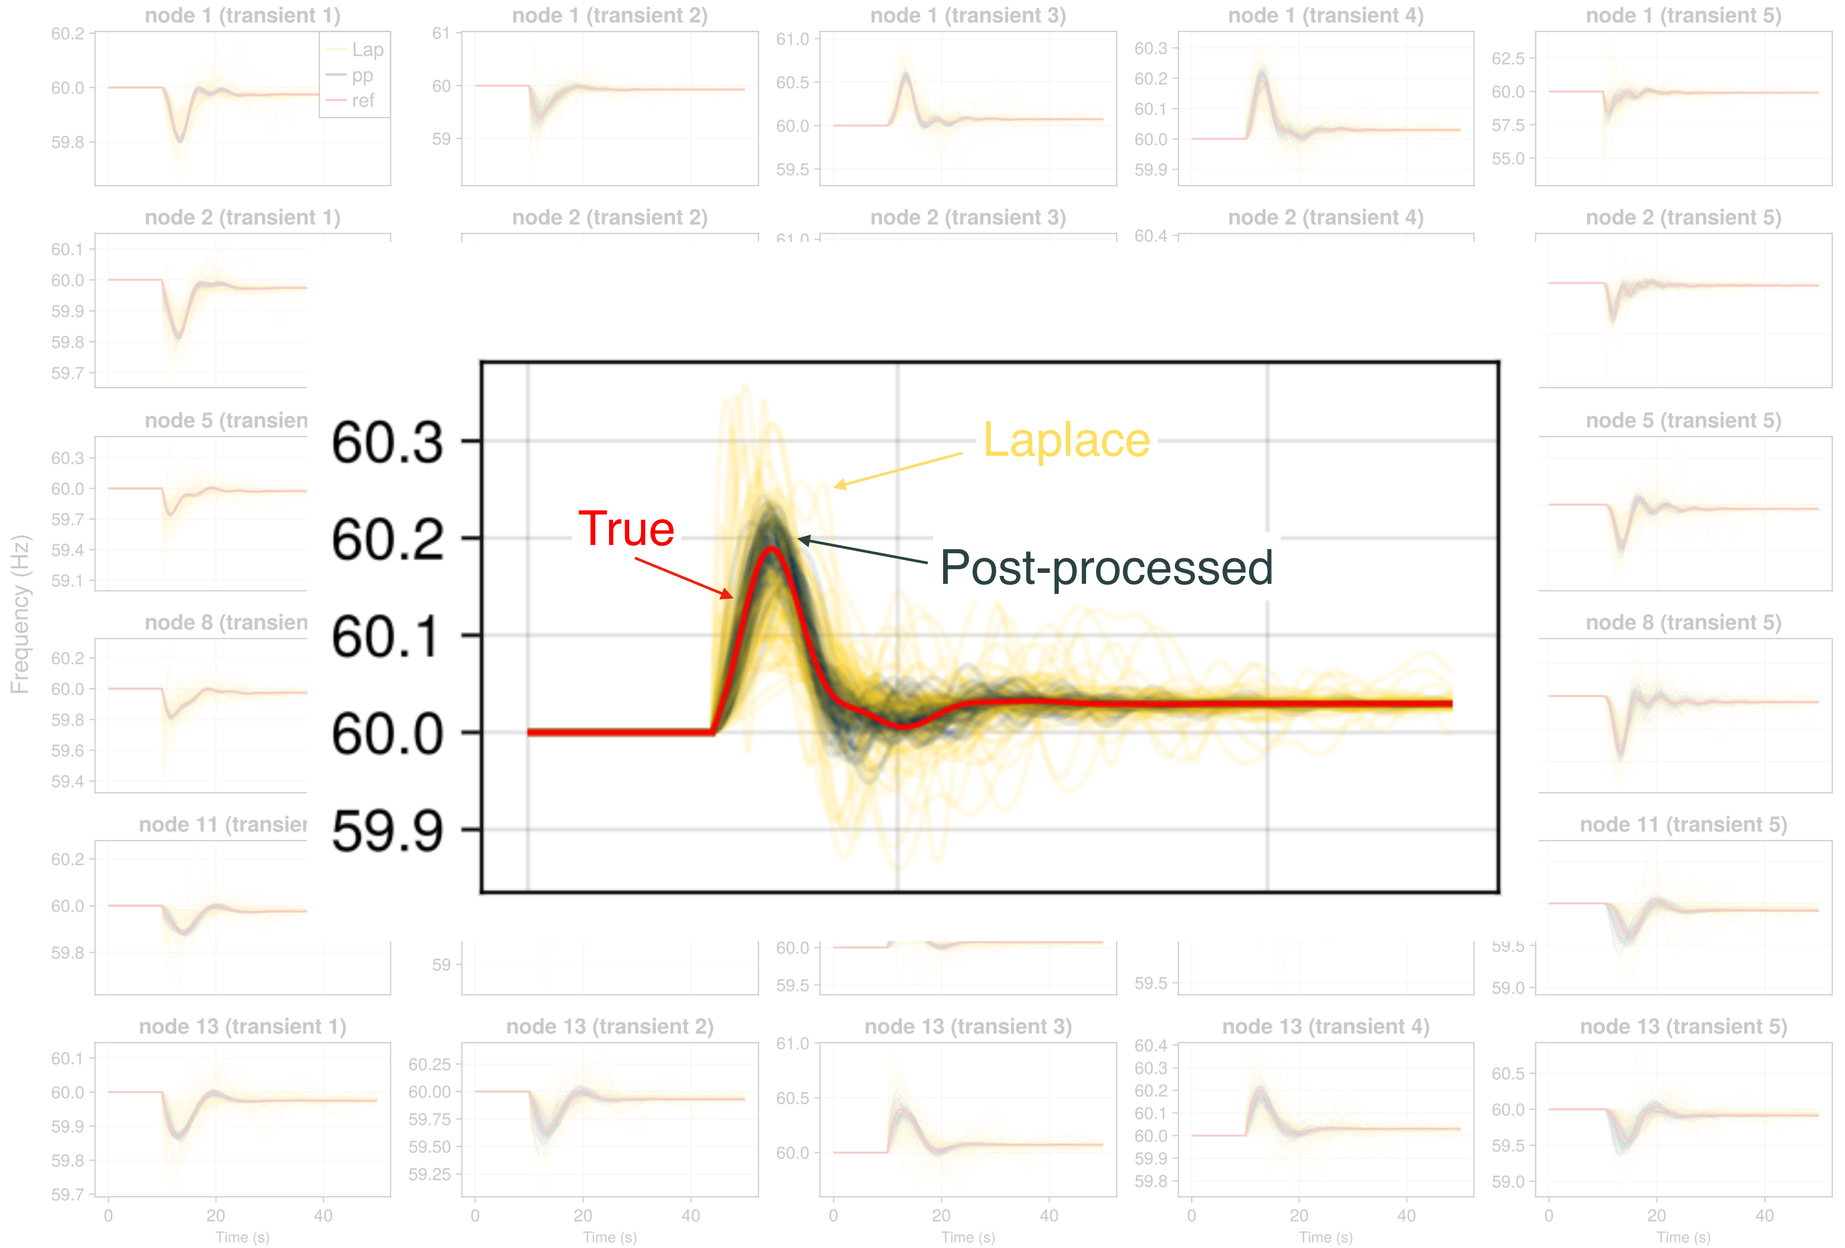

Laplace mechanism vs. post-processing: Model fidelity and resulting dynamics

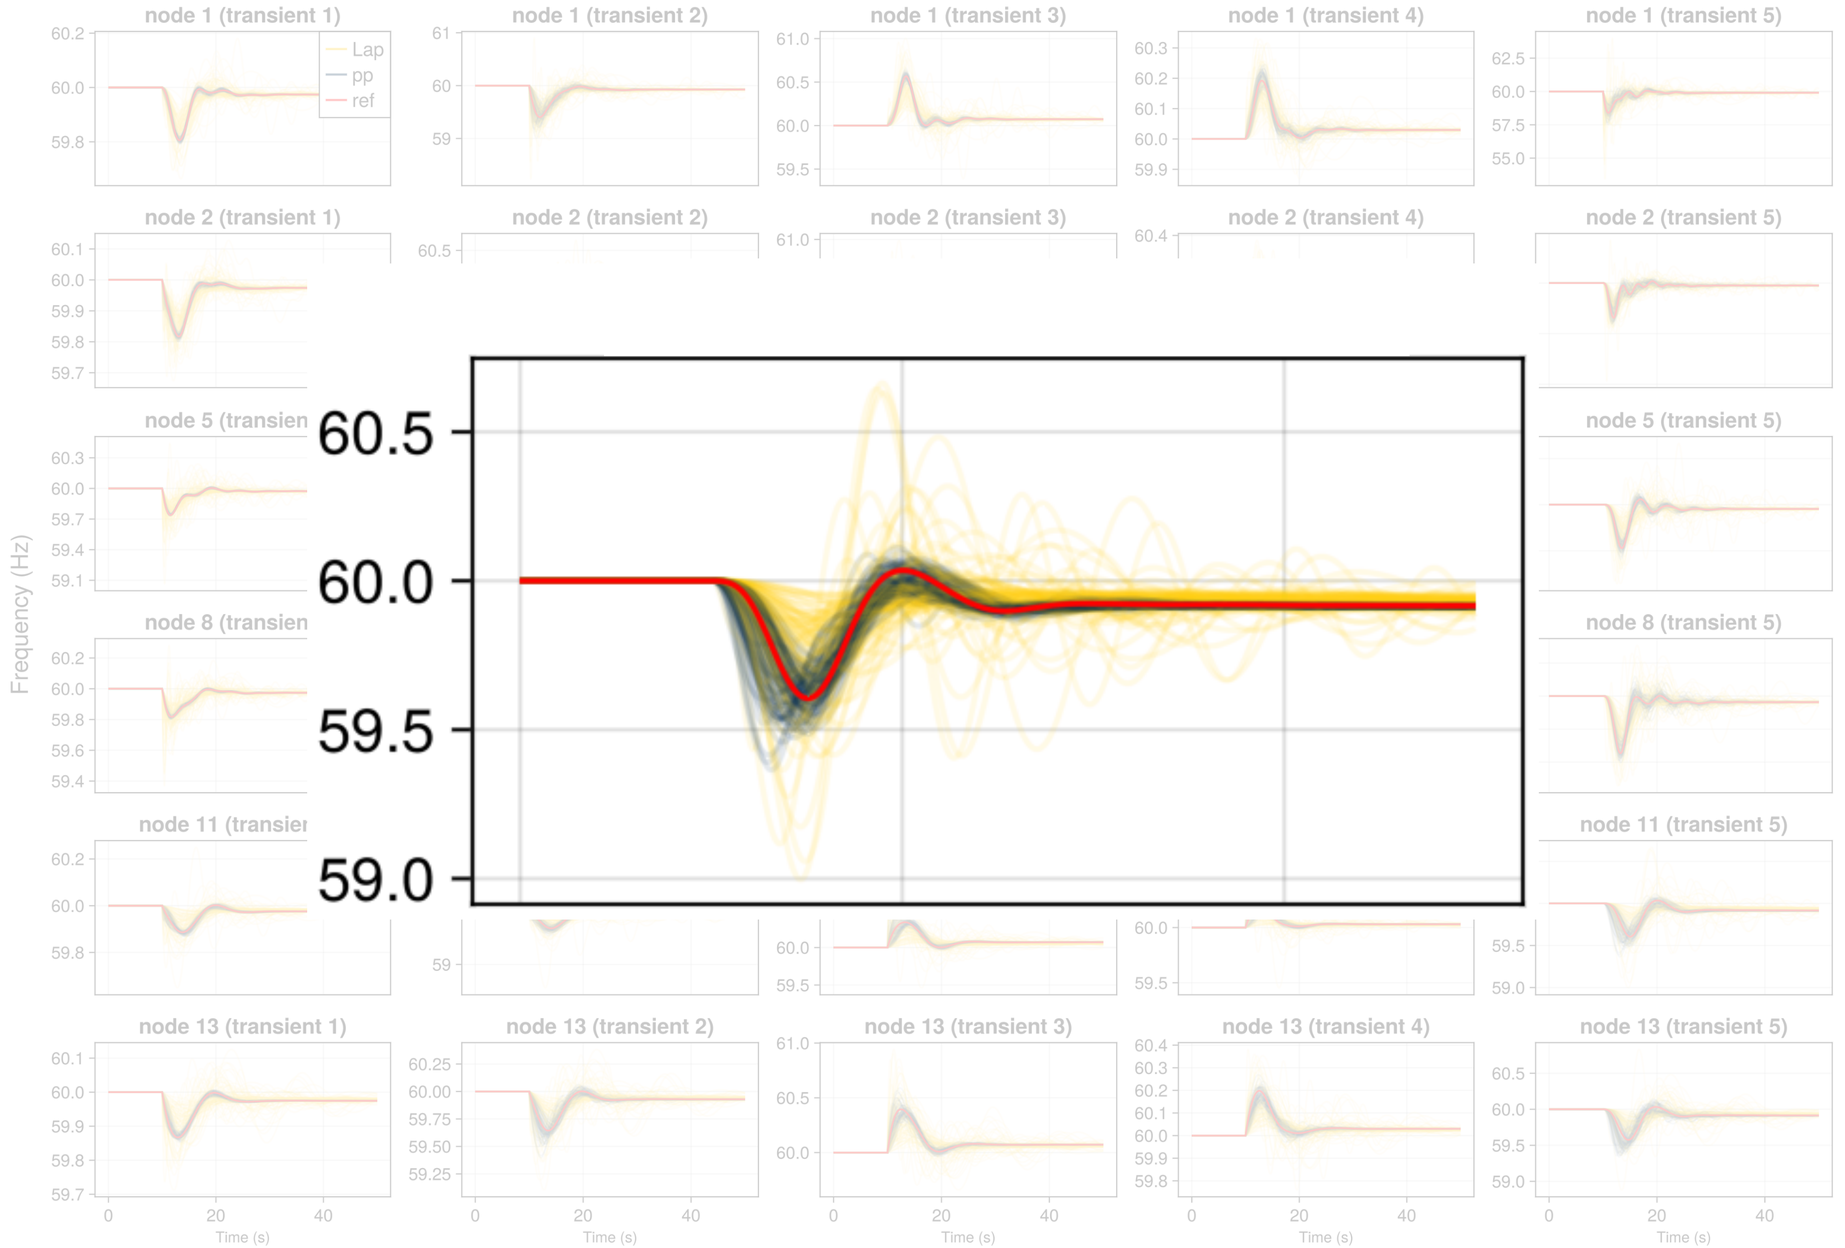

Agenda

1. Intro
2. Formalization of energy data privacy
3. Synthesizing optimization data with privacy and cyber security guarantees
4. Synthesizing power system dynamics models with privacy guarantees
5. **Outro**

What we have in other fields….

Computer vision
(ImageNet)

Speech recognition
(LibriSpeech)

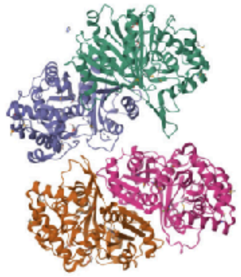

Biology
(UniProt)

What we *can* have in power systems

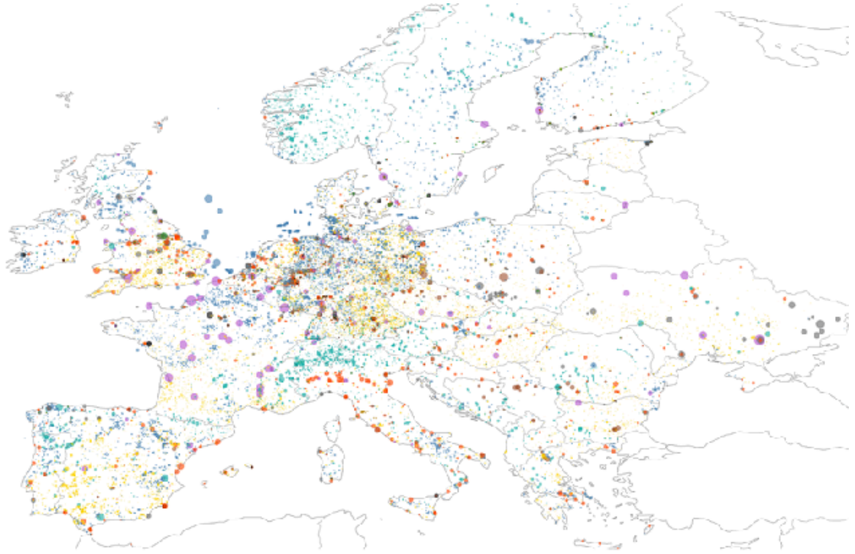

- Statistically credible and operation-feasible models of power grid dispatch
- High-fidelity models of power systems dynamics
- Arbitrary large, credible training datasets for machine learning applications in power systems

Future of synthetic power system datasets

What we used to say about synthetic datasets:

- "[...] data bears **no relation** to the actual grid [...]"
- "This test case represents [...] **fictitious** transmission"
- "This case is synthetic and **does not** model the actual grid"

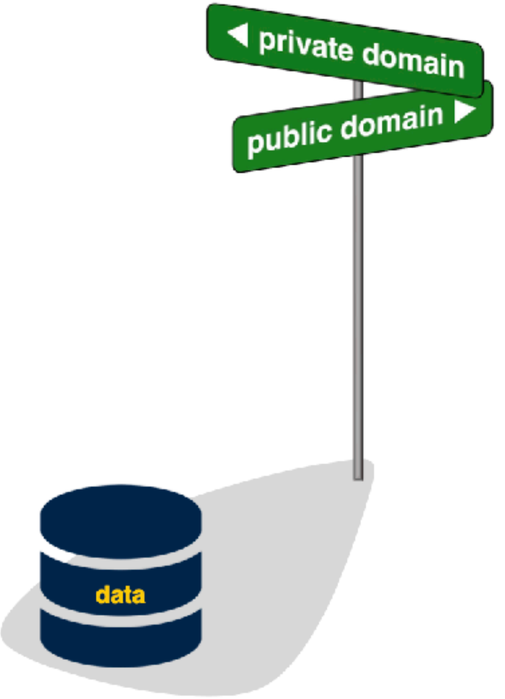

What we will say about synthetic datasets:

- "This synthetic dataset is produced based on the data from a real-world power grid"
- "It is not possible to infer the real data from this synthetic dataset"
- "Computational results on this data are consistent with the real data"

From today's talk

Synthetic optimization data

438 IEEE CONTROL SYSTEMS LETTERS, VOL. 9, 2025

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.**

I. Introduction

OPTIMAL power flow (OPF) analysis in power systems requires realistic grid models with accurate network, generation, and load parameters—data that is difficult to source from real-world grids due to privacy and (cyber-)security concerns. While the lack of such models has inspired the development of artificial grids [1], [2], a more principled approach leverages the theory of differential privacy (DP) [3] to release grid models directly from real-world systems.

The DP theory asserts that it is impossible—up to prescribed privacy parameters—to infer the original parameters from their DP release. Such strong privacy guarantees originate from Laplacian perturbations [4] of real grid parameters, followed by post-processing optimization of the perturbed parameters to restore their modeling fidelity to the source grid, e.g., in terms of similarity of the OPF outcomes [5], [6], [7]. The DP theory also lies at the core of modern privacy-preserving OPF solvers [8], [9], [10], the release of aggregated grid statistics [11], and related grid information [12].

However, the privacy guarantees alone may not suffice to release grid parameters, as cybersecurity risks associated 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 build on [5], [6], [7] and leverage the Laplace mechanism and post-processing optimization to tune synthetic data while anticipating cyber risks through embedded attack optimization.

The contributions of this letter are summarized as follows:

1) We formulate a Cyber Resilient Obfuscation (CRO) algorithm, an optimization-based algorithm to release electric load data with a guarantee to preserve the privacy of the original data and ensure the resilience of the source grid to load redistribution attacks. The algorithm post-processes synthetic loads to balance their fidelity with the potential damage to the grid; other grid

Received 16 March 2025; revised 2 May 2025; accepted 15 May 2025. Date of publication 27 May 2025; date of current version 11 June 2025. Recommended by Senior Editor S. Olaru. *(Corresponding author: Shengyang Wu.)*

The authors are with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109 USA (e-mail: syseanwu@umich.edu; dvorkin@umich.edu).

Digital Object Identifier 10.1109/LCSYS.2025.3574146

Synthetic dynamic models

Synthetic machine learning datasets

Constrained Diffusion Models for Synthesizing Representative Power Flow Datasets

Milad Hoseinpour, *Student Member, IEEE*, Vladimir Dvorkin, *Member, IEEE*

***Abstract*—High-quality power flow datasets are essential for training machine learning models in power systems. However, security and privacy concerns restrict access to real-world data, making statistically accurate and physically consistent synthetic datasets a viable alternative. We develop a diffusion model for generating synthetic power flow datasets from real-world power grids that both replicate the statistical properties of the real-world data and ensure AC power flow feasibility. To enforce the constraints, we incorporate gradient guidance based on the power flow constraints to steer diffusion sampling toward feasible samples. For computational efficiency, we further leverage insights from the fast decoupled power flow method and propose a variable decoupling strategy for the training and sampling of the diffusion model. These solutions lead to a physics-informed diffusion model, generating power flow datasets that outperform those from the standard diffusion in terms of feasibility and statistical similarity, as shown in experiments across IEEE benchmark systems.**

***Index Terms*—Diffusion model, generative AI in power systems, physics-informed machine learning, power flow, synthetic data.**

I. Introduction

POWER flow datasets [1]–[3] are essential for training and benchmarking machine learning (ML) models for optimal power flow (OPF) [4] and state estimation [5]. However, the real-world power flow datasets are rarely available due to privacy, security, and legal barriers [6]–[10]. Recent advances in generative AI, capable of producing synthetic data with distributions similar to the original data [11]–[20], have partially lifted these barriers, yet statistical consistency alone cannot guarantee adherence to physical grid constraints [21]. Consequently, ML models trained on constraint-agnostic synthetic datasets are likely to perform substantially worse than those trained on original data. This paper introduces a data generation framework to synthesize statistically consistent and physically meaningful power flow datasets. To achieve this, we develop a constrained diffusion model to learn the underlying distribution of power flow data and generate synthetic samples that are both statistically representative and feasible with respect to the AC power flow constraints. This constrained diffusion model can be trained internally by system operators to publicly release high-quality synthetic power flow data to support a wide range of downstream ML applications.

A. Related Work

The literature on generating synthetic datasets for power systems broadly falls into two categories: generic random sampling and historical data-driven approaches.

The former focuses on power flow data generation through iterative uniform sampling of loads followed by solving the OPF problem [22], [23]. In [24], authors use a truncated Gaussian distribution as another variation of sampling, which also accounts for correlations between power injections at different locations. However, the datasets based on generic sampling only represent a small portion of the feasibility region. To solve this, [6] uniformly samples loads from a convex set, containing the feasible region, and iteratively refines this set using infeasibility certificates. In [25], a bilevel optimization is proposed to sample operating conditions close to the boundaries of the feasible region, which is more informative that a random sampling. A basic requirement for ML-based OPF solvers is robustness to grid topology variations, e.g., network topology switching [26]. To meet this requirement, authors in [27] incorporate topological perturbations in addition to load perturbations in their synthetic data generation framework.

Although straightforward, random sampling comes with certain limitations. The resulting datasets do not represent the true underlying distribution of real-world operating conditions. That is, the synthetic data points may fail to capture correlations, patterns, or variability present in historical data. ML-based OPF solvers trained on such data may generalize poorly, leading to inaccurate predictions and erroneous uncertainty quantification [28], [29]. Moreover, the required number of random samples to cover the whole feasible region grows ex

The authors are with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109, USA. E-mail: {miladh, dvorkin}@umich.edu.

(Related work on diffusion for power flows)

Thank you for your attention!