High-dimensional statistics for electric power systems

Samuel Talkington

Ph.D. Student, Georgia Tech ECE

AI4OPT 3MT Competition

April 1, 2025

My background

My great-grandfather, an electrician, helped electrify Appalachia.

I'm a native of central West Virginia, a cornerstone of our nation's energy infrastructure.


MONONGAHELA NEWS
ON THE JOB

Singled out for special mention when the 132 kv metering equipment was cut in recently at the National Carbon Company's Clarksburg plant was Johnny Singleton, shown here beside one of the giant transformers on location. Johnny was on the job throughout the installation and knows the operation from A to Izzard.

1950

Motivation: What node in a network is most probable to have a constraint violation?

...Under mild assumptions, and ***given realistic parameters****...*

Previous work

Traditional sensor placement: offline, static

With sensors

Without sensors

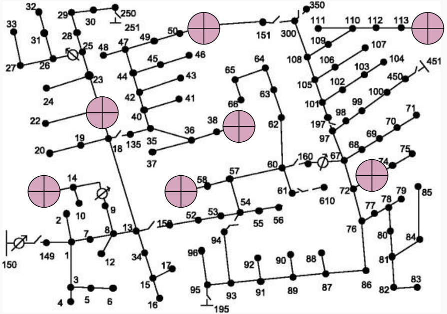

Select *locations* (a subset of nodes) to install sensors and measure continuously

Motivation

Dave Rieken,
Vice President of
Research

Sensor *placement* in power systems is **wrong**!

We have sensors at ***every* node:** The problem is ***how*** **we sense!**

We make smart meters, folks!

70% of North American households have smart meters, EIA, 2020.

Here's an algorithm that selects sensor locations

Samuel Talkington

Sensor placement–or sampling? Rethinking sensor placement.

Dave Rieken,
Vice President of Research

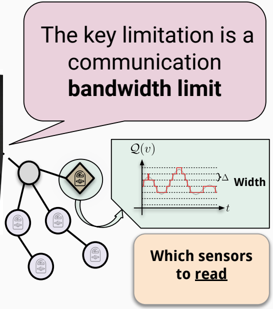

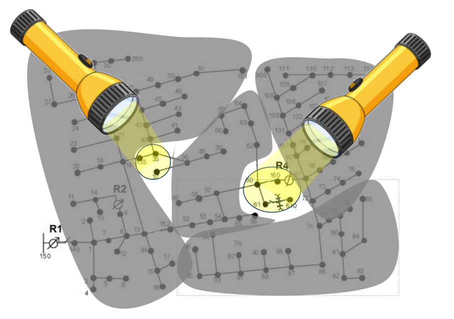

70% of North American households have smart meters, EIA, 2020.

The problem

- Power distribution networks have high levels of sensors already, but with…
- *Limited communication **bandwidth**.*
- **How do we dynamically monitor these sensor networks efficiently?**
- **i.e., how to move these flashlights around?**

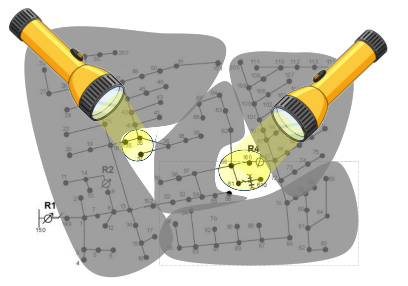

Select only a few sensors

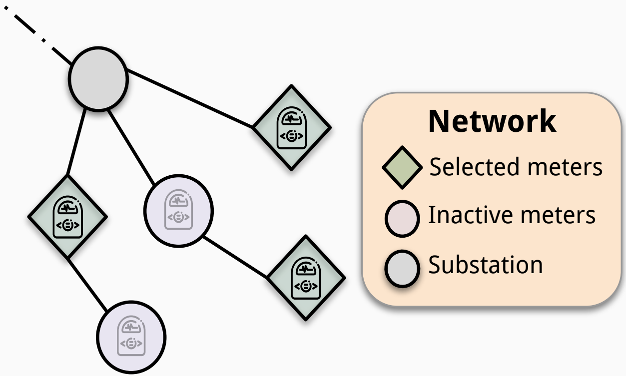

Figure 4: **Key idea:** we can only select a few sensors

Select $\mathcal{S}$, find worst case in $\mathcal{S}$

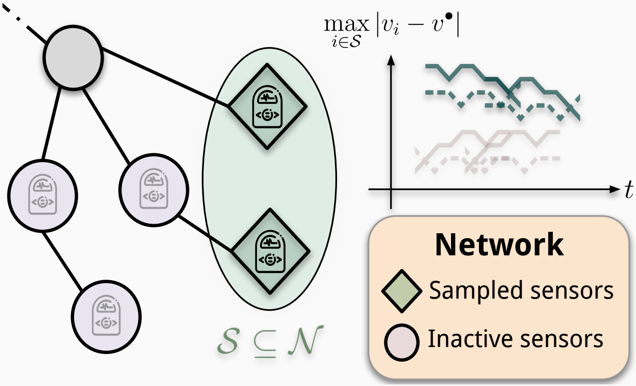

Figure 5: From $\mathcal{S}$, what's the **worst case** voltage?

Strategy

At each time t: the learner picks b nodes to **check the security.**

The set of all strategies is the **set of all subsets of b nodes.**

$$\mathcal{A} = \left\{ \mathcal{S} \in 2^{\mathcal{N}} \; : \; |\mathcal{S}| \leq b \right\}, \tag{7}$$

so there are $|\mathcal{A}| = \binom{n}{b}$ possible strategies…challenging in general!

Reward

When the learner has selected sensors $\mathcal{S}_t \in \mathcal{A}$ to *ping*, she observes a *reward* $f : \mathcal{A} \to \mathbb{R}$ that looks like

$$\boxed{f(\mathcal{S}) = \text{Worst case voltage in } \mathcal{S}.} \tag{8}$$

In symbols:

$$f(\mathcal{S}) = \max_{i \in \mathcal{S}_t} |\epsilon_i| = \max_{i \in \mathcal{S}_t} |v_i - 1| = \max_{i \in \mathcal{S}_t} |\langle \mathbf{w}_i, \psi \rangle| \,. \tag{9}$$

This reward is the **maximum voltage magnitude observed in the sampling strategy.**

How to optimize sensor sampling: A bandit approach

To pick the best *m*-sample strategy, minimize the **regret**:

$$\boxed{R_m := \mathrm{E}\left[\text{Best voltage sampling strategy} - \text{Your voltage sampling strategy}\right]}$$

Intuition: Iteratively track extreme voltages based on feedback.

Spectral bandit algorithm

Solution approach: At each timestep t, recursively compute an estimate of the *Fourier coefficients* ψ for the voltage magnitudes $\boldsymbol{v}$:

$$\hat{\boldsymbol{\psi}}_t = \underset{\boldsymbol{\psi}\in\mathbb{R}^n}{\arg\min} \sum_{s=1}^{t-1} (v_s - \langle \boldsymbol{w}_s, \boldsymbol{\psi} \rangle)^2 + \beta \, ||\boldsymbol{\psi}||^2_{\boldsymbol{\Lambda}} \,, \tag{10}$$

where $\beta > 0$ is a regularization parameter that you choose. The indices $s = 1, \dots t-1$ are the **sampled nodes!**

Spectral regularization

The regularization term, $||\psi||_{\mathbf{\Lambda}}$, promotes predictions of the voltages that are **electrically diverse**:

$$||\psi||_{\mathbf{\Lambda}} := \sqrt{\psi^{\mathsf{T}}\mathbf{\Lambda}\psi} = \sqrt{\sum_{(i,j)\in\mathcal{E}} y_{ij}(\psi_i - \psi_j)^2}. \tag{11}$$

This is also known as the **Dirichlet energy** of the graph.

Relates to *effective resistance*...check out the paper for more information

Intuition of spectral regularization

Spectral regularization promotes an **electrically diverse** sampling policy

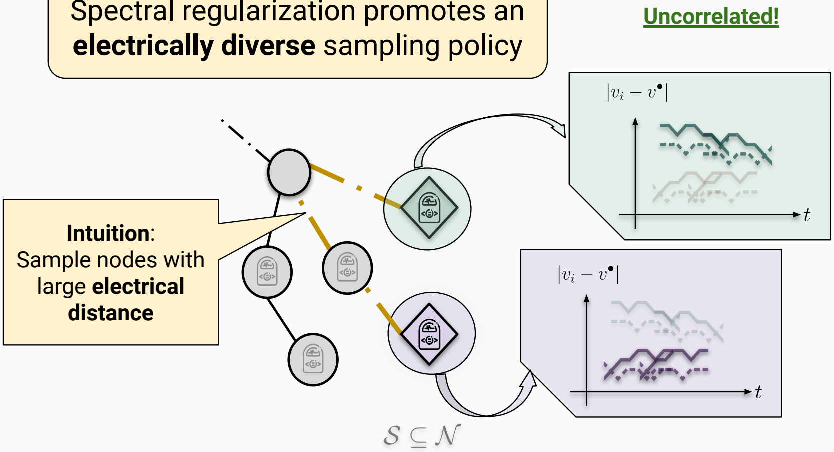

Key take-away

Question: Why is this an improvement?

Answer: The worst case m-sample regret with conventional linear bandits is

$$R_m \leq \tilde{O}(n\sqrt{m}),$$

where n is the number of nodes. Our result, by **incorporating the graphical structure**[1] of the power flow equations, yields

$$R_m \leq \tilde{O}(d\sqrt{m}).$$

This reduces the scaling factor to the *intrinsic dimension*, $d < n$, of the graph Laplacian.

(This is a huge improvement, as we will see empirically.)

[1] T. Kocák, et al., "Spectral Bandits", Journal of Machine Learning Research, 21 (1), Jan. 2020.

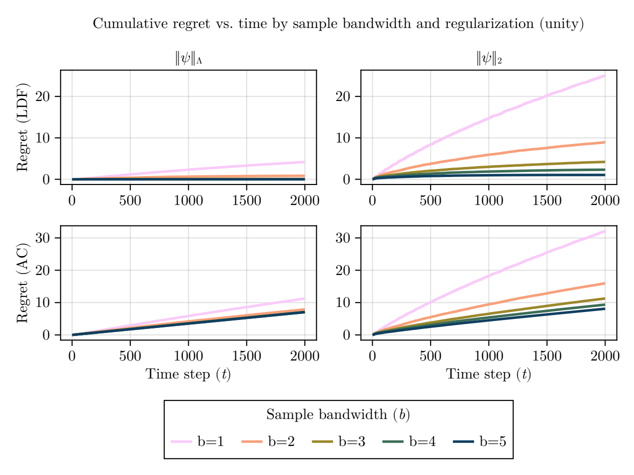

Figure 6: Fixed power factor: Regret of the bandwidth-constrained maximal voltage risk sampler vs. time with spectral (left) and ℓ_2 (right) regularization.

Thanks! Keep in touch: talkington@gatech.edu

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-2039655. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the authors(s) and do not necessarily reflect the views of the National Science Foundation.