Variability in power systems: stochastic defense against ideal grid attacks

Daniel Bienstock, Mauro Escobar, Apurv Shukla

Columbia University

Mopta 2018

Variability in power systems

- The next frontier: controlling short-term variability (seconds or less)

Variability in power systems

- The next frontier: controlling short-term variability (seconds or less)
- Goal: safety and controllability as much as economics, or more

Variability in power systems

- The next frontier: controlling short-term variability (seconds or less)
- Goal: safety and controllability as much as economics, or more
- Driven by smart loads, DERs, DPVs, solid state devices, batteries, etc ... and

Variability in power systems

- The next frontier: controlling short-term variability (seconds or less)
- Goal: safety and controllability as much as economics, or more
- Driven by smart loads, DERs, DPVs, solid state devices, batteries, etc ... and **PMUs**

Variability in power systems

- The next frontier: controlling short-term variability (seconds or less)
- Goal: safety and controllability as much as economics, or more
- Driven by smart loads, DERs, DPVs, solid state devices, batteries, etc ... and **PMUs**
- PMUs = "phasor measurement units," relatively expensive but the way of the future

Variability in power systems

- The next frontier: controlling short-term variability (seconds or less)
- Goal: safety and controllability as much as economics, or more
- Driven by smart loads, DERs, DPVs, solid state devices, batteries, etc ... and **PMUs**
- PMUs = "phasor measurement units," relatively expensive but the way of the future
- Goal: very tight, near-real-time control of power systems

Variability in power systems

- The next frontier: controlling short-term variability (seconds or less)
- Goal: safety and controllability as much as economics, or more
- Driven by smart loads, DERs, DPVs, solid state devices, batteries, etc ... and **PMUs**
- PMUs = "phasor measurement units," relatively expensive but the way of the future
- Goal: very tight, near-real-time control of power systems
- Must be able to learn real-time structure and stochastics

Variability in power systems

- The next frontier: controlling short-term variability (seconds or less)
- Goal: safety and controllability as much as economics, or more
- Driven by smart loads, DERs, DPVs, solid state devices, batteries, etc ... and **PMUs**
- PMUs = "phasor measurement units," relatively expensive but the way of the future
- Goal: very tight, near-real-time control of power systems
- Must be able to learn real-time structure and stochastics
- Joint work: Columbia and LANL

"Cyber-physical" attacks on power grids

"Cyber-physical" attacks on power grids

Fact or fiction?

1. An adversary carries out a physical alteration of a grid (example: disconnecting a power line)

"Cyber-physical" attacks on power grids

Fact or fiction?

1. An adversary carries out a physical alteration of a grid (example: disconnecting a power line)
2. This is coordinated with a modification of sensor signals – a **hack**

"Cyber-physical" attacks on power grids

Fact or fiction?

1. An adversary carries out a physical alteration of a grid (example: disconnecting a power line)
2. This is coordinated with a modification of sensor signals – a **hack**
3. The goal is to disguise, or keep completely hidden, the nature of the attack and its likely consequences

"Cyber-physical" attacks on power grids

Fact or fiction?

1. An adversary carries out a physical alteration of a grid (example: disconnecting a power line)
2. This is coordinated with a modification of sensor signals – a **hack**
3. The goal is to disguise, or keep completely hidden, the nature of the attack and its likely consequences
4. Power industry: it will never happen
   ("we would know what happened")

“Cyber-physical” attacks on power grids

Fact or fiction?

1. An adversary carries out a physical alteration of a grid (example: disconnecting a power line)
2. This is coordinated with a modification of sensor signals – a **hack**
3. The goal is to disguise, or keep completely hidden, the nature of the attack and its likely consequences
4. Power industry: it will never happen
   (“we would know what happened”)
5. Really?

Control centers, RTUs, PMUs, state estimation

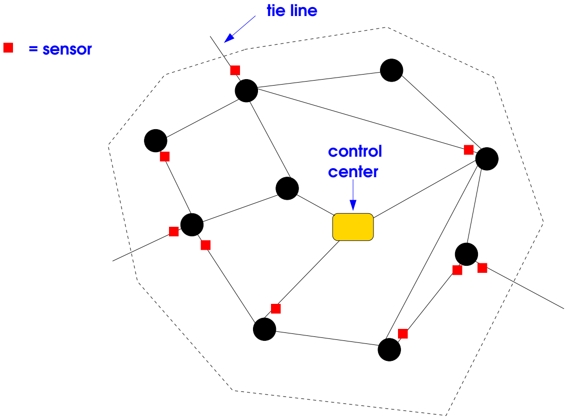

Control centers, RTUs, PMUs

- Control center performs a regulatory and economic role
- Sensors report to control center
- Control center issues commands to (in particular) smaller generators
- Sensors: RTUs (old), PMUs (new – and more expensive)
- RTUs report **once every four seconds**
- PMUs report
  - **30 to 100 times a second**
  - PMUs report (AC) voltage and current (plus more ...)
- Anecdotal: PMUs overwhelming human operators
- But PMUs are the way of the future

State estimation (very abridged)

A data-driven procedure to estimate relevant grid parameters

- Even with PMUs, data can be "complex"
- Statistical procedure: "state estimation" (at control center)

DC power flow equations:

$$B\theta \ = \ P^g - P^d$$

$\boldsymbol{B}$ = susceptance matrix, $\boldsymbol{\theta}$ = phase angles, $\boldsymbol{P^g}$, $\boldsymbol{P^d}$ generation and load vectors

- Sensors provide information that fit **some** of the $\boldsymbol{\theta}, \boldsymbol{P^d}, (\boldsymbol{P^g}?)$ parameters
- State estimation: least squares procedure to estimate the rest, plus more

Some prior basic research on “cyberphysical” attacks

- Intelligent procedures for enriching state estimation so as to detect and reconstruct attacks
- Unavoidable: a model for attacking behavior is essential
- Liu Ning Reiter (2009), Kim Poor (2011),
- Deka Baldick Vishwanath (2015)
- Soltan Yannakakis Zussman (2015 -)
- **Warning:** watch out for those assumptions!
- Attacks are **static**

Some prior basic research on "cyberphysical" attacks

- Intelligent procedures for enriching state estimation so as to detect and reconstruct attacks
- Unavoidable: a model for attacking behavior is essential
- Liu Ning Reiter (2009), Kim Poor (2011),
- Deka Baldick Vishwanath (2015)
- Soltan Yannakakis Zussman (2015 -)
- **Warning:** watch out for those assumptions!
- Attacks are **static** and defense is **passive**

Today: load change, signal hacking – all AC

- An attacker causes physical changes in the network:

Today: load change, signal hacking – all AC

- An attacker causes physical changes in the network:
- In particular **load** changes (no generator changes)
- Possibly also line disconnections

Today: load change, signal hacking – all AC

- An attacker causes physical changes in the network:
- In particular **load** changes (no generator changes)
- Possibly also line disconnections
- Attacker also hacks the signal flow: the output of some sensors is altered
- Goal of the attacker is twofold:
  1. Hide the location of the attack and even the fact that an attack happened
  2. **Cause line overloads that remain hidden**

Today: load change, signal hacking – all AC

- An attacker causes physical changes in the network:
- In particular **load** changes (no generator changes)
- Possibly also line disconnections
- Attacker also hacks the signal flow: the output of some sensors is altered
- Goal of the attacker is twofold:
  1. Hide the location of the attack and even the fact that an attack happened
  2. **Cause line overloads that remain hidden**
- Attacker expects **full PMU deplyoment**. Everything is **AC** based.

Basic AC model of a power line in steady state

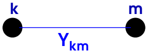

- Line between buses (nodes) $\boldsymbol{k}$ and $\boldsymbol{m}$.

- $\boldsymbol{Y_{km}}$: 2×2 (complex) admittance matrix (physics of the line)
- $\boldsymbol{V_k}$ = voltage at k $= |\boldsymbol{V_k}|\boldsymbol{e^{j\theta_k}}$, $\boldsymbol{j = \sqrt{-1}}$, similarly with $\boldsymbol{V_m}$

- Current-voltage relationship:

$$\begin{pmatrix} I_{km} \\ I_{mk} \end{pmatrix} = Y_{km} \begin{pmatrix} V_k \\ V_m \end{pmatrix}$$

- $\boldsymbol{I_{km}, I_{mk}}$ = (complex) current injected into line at $\boldsymbol{k}$ (resp. $\boldsymbol{m}$)

- $\boldsymbol{S_{km}}$ = (complex) power injected into line at $\boldsymbol{k}$ = $\boldsymbol{V_k I^*_{km}}$

What happens when there is a generation/load mismatch

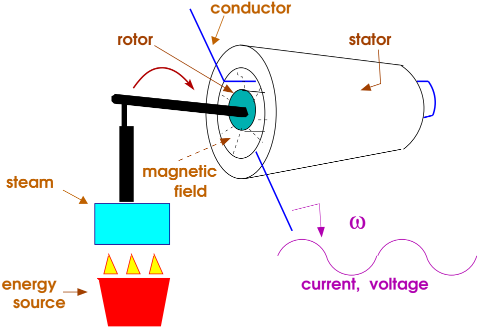

What happens when there is a generation/load mismatch

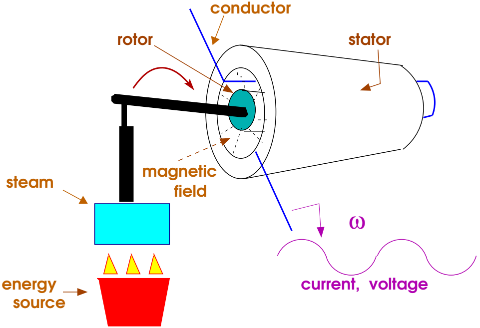

Frequency response:

What happens when there is a generation/load mismatch

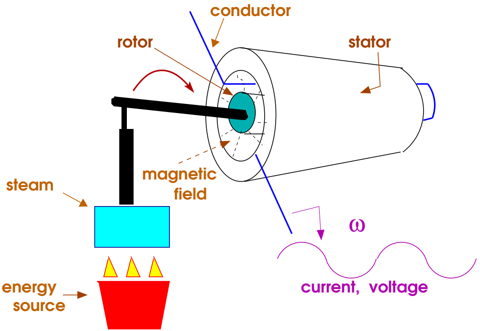

Frequency response:

mismatch $\boldsymbol{\Delta P}$

What happens when there is a generation/load mismatch

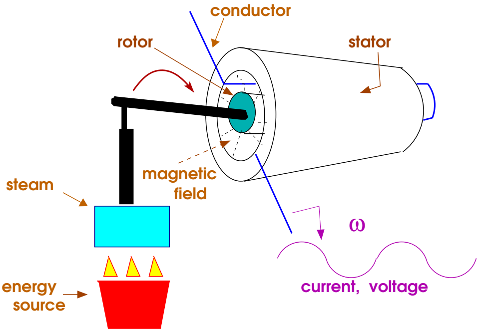

Frequency response:

mismatch $\boldsymbol{\Delta P}$ $\Rightarrow$ **frequency change** $\boldsymbol{\Delta\omega} \approx -c\,\Delta P$

AGC, primary and secondary response (simplified!, abridged!)

Suppose generation vs loads balance spontaneously changes (i.e. a net imbalance)?

- AC frequency changes proportionally (to first order) near-instantaneously

AGC, primary and secondary response (simplified!, abridged!)

Suppose generation vs loads balance spontaneously changes (i.e. a net imbalance)?

- AC frequency changes proportionally (to first order) near-instantaneously
- **Primary response.** (very quick) Inertia in generators contributes electrical energy to the system

AGC, primary and secondary response (simplified!, abridged!)

Suppose generation vs loads balance spontaneously changes (i.e. a net imbalance)?

- AC frequency changes proportionally (to first order) near-instantaneously
- **Primary response.** (very quick) Inertia in generators contributes electrical energy to the system
- **Secondary response.** (seconds) Suppose **estimated** generation **shortfall** = $\boldsymbol{\Delta P}$. Then:

AGC, primary and secondary response (simplified!, abridged!)

Suppose generation vs loads balance spontaneously changes (i.e. a net imbalance)?

- AC frequency changes proportionally (to first order) near-instantaneously
- **Primary response.** (very quick) Inertia in generators contributes electrical energy to the system
- **Secondary response.** (seconds) Suppose **estimated** generation **shortfall** $= \boldsymbol{\Delta P}$. Then:

Generator $\boldsymbol{g}$ changes output by $\alpha_{\boldsymbol{g}} \boldsymbol{\Delta P}$

AGC, primary and secondary response (simplified!, abridged!)

Suppose generation vs loads balance spontaneously changes (i.e. a net imbalance)?

- AC frequency changes proportionally (to first order) near-instantaneously
- **Primary response.** (very quick) Inertia in generators contributes electrical energy to the system
- **Secondary response.** (seconds) Suppose **estimated** generation **shortfall** = $\boldsymbol{\Delta P}$. Then:

Generator $\boldsymbol{g}$ changes output by $\boldsymbol{\alpha_g \Delta P}$

- $\sum_g \alpha_g = 1$, $\alpha \geq 0$,

AGC, primary and secondary response (simplified!, abridged!)

Suppose generation vs loads balance spontaneously changes (i.e. a net imbalance)?

- AC frequency changes proportionally (to first order) near-instantaneously
- **Primary response.** (very quick) Inertia in generators contributes electrical energy to the system
- **Secondary response.** (seconds) Suppose **estimated** generation **shortfall** $= \boldsymbol{\Delta P}$. Then:

Generator $\boldsymbol{g}$ changes output by $\boldsymbol{\alpha_g \Delta P}$

- $\sum_g \alpha_g = 1$, $\alpha \geq 0$, $\alpha > 0$ for "participating" generators
- **Preset** participation factors
- $\boldsymbol{\Delta \omega}$ sensed by control center, which issues generator commands

Ideal ("perfect") static attack: setup

- PMUs everywhere: at both ends of each line

Ideal ("perfect") static attack: setup

- PMUs everywhere: at both ends of each line
- Attacker has been in the system long enough to learn the system (data-wise)

Ideal ("perfect") static attack: setup

- PMUs everywhere: at both ends of each line
- Attacker has been in the system long enough to learn the system (data-wise)
- Attacker chooses, in advance, a non-generator, sparse set $\mathcal{A}$ of buses to attack and in particular a line uv to overload

Ideal ("perfect") static attack: setup

- PMUs everywhere: at both ends of each line
- Attacker has been in the system long enough to learn the system (data-wise)
- Attacker chooses, in advance, a non-generator, sparse set $\mathcal{A}$ of buses to attack and in particular a line uv to overload
- In near real-time, the attacker learns the current loads, up to small error

Ideal ("perfect") static attack: setup

- PMUs everywhere: at both ends of each line
- Attacker has been in the system long enough to learn the system (data-wise)
- Attacker chooses, in advance, a non-generator, sparse set $\mathcal{A}$ of buses to attack and in particular a line uv to overload
- In near real-time, the attacker learns the current loads, up to small error
- In near real-time, the attacker solves computational problem that diagrams the attack on $\mathcal{A}$

Ideal ("perfect") static attack: setup

- PMUs everywhere: at both ends of each line
- Attacker has been in the system long enough to learn the system (data-wise)
- Attacker chooses, in advance, a non-generator, sparse set $\mathcal{A}$ of buses to attack and in particular a line uv to overload
- In near real-time, the attacker learns the current loads, up to small error
- In near real-time, the attacker solves computational problem that diagrams the attack on $\mathcal{A}$
- This will specify the load changes and the signal distortion
- Post-attack, attacker cannot recompute much and only relies on adding "noise" to the computed distorted signals

Undetectable attack: The attacker's perspective

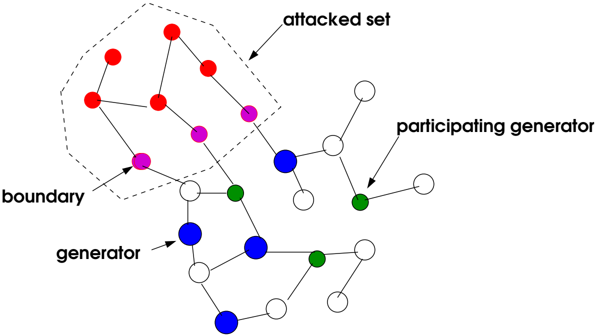

Undetectable attack: decisions for the attacker (abridged!)

- For every bus in $\mathcal{A}$, a "true" and "reported" complex voltage

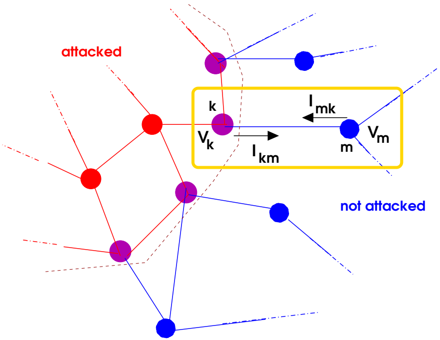

Undetectable attack: decisions for the attacker (abridged!)

- For every bus in $\mathcal{A}$, a "true" and "reported" complex voltage

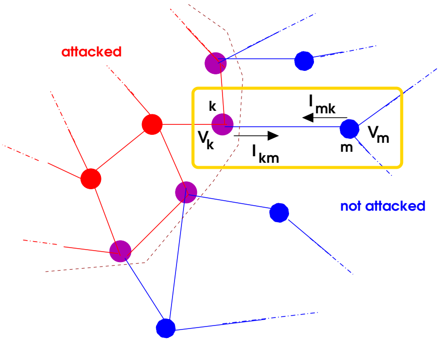

But: $\begin{pmatrix} I_{km} \\ I_{mk} \end{pmatrix} = Y_{km} \begin{pmatrix} V_k \\ V_m \end{pmatrix}$ **so** V_k must be truthful

Undetectable attack: tasks for the attacker (abridged!)

- For every bus in $\mathcal{A}$, compute a "true" and "reported" complex voltage (magnitude and angle) V_k^{T} and V_k^{R}
- True and reported voltages **must** agree on the boundary of $\mathcal{A}$!
- Compute true and reported **currents** for lines within $\mathcal{A}$
- Compute voltages and currents on all other lines (true and reported are identical)
- Compute two power flow solutions; each must satisfy AC power equations, load changes a variable
- On responding generators: compute generation change consistent with secondary response if loads are modified
- Restriction to attacker: attacked zone does not include **any** generators.

Undetectable attack: tasks for the attacker (abridged!)

- For every bus in $\mathcal{A}$, compute a "true" and "reported" complex voltage (magnitude and angle) V_k^{T} and V_k^{R}
- True and reported voltages **must** agree on the boundary of $\mathcal{A}$!
- Compute true and reported **currents** for lines within $\mathcal{A}$
- Compute voltages and currents on all other lines (true and reported are identical)
- Compute two power flow solutions; each must satisfy AC power equations, load changes a variable
- On responding generators: compute generation change consistent with secondary response if loads are modified
- Restriction to attacker: attacked zone does not include **any** generators. Why?
- Some additional lying

Undetectable static attack (load modification, no line tripping, abridged!)

$$\textbf{Max} \quad (p_{uv}^{\mathrm{T}})^2 + (q_{uv}^{\mathrm{T}})^2 \quad \textbf{square norm of flow on uv} \tag{1a}$$

s.t.

$$\forall k \in \mathcal{A}^C \cup \partial\mathcal{A}, \; V_k^{\mathrm{T}} = V_k^{\mathrm{R}} \; \textbf{(truthful voltages outside attacked zone)} \tag{1b}$$

$$\forall k \in \mathcal{A}, \; -(P_k^{d,\mathrm{R}} + jQ_k^{d,\mathrm{R}}) = \sum_{km \in \delta(k)} (p_{km}^{\mathrm{R}} + jq_{km}^{\mathrm{R}}), \; \textbf{(true power flow balance in attacked zone)} \tag{1c}$$

$$-(P_k^{d,\mathrm{T}} + jQ_k^{d,\mathrm{T}}) = \sum_{km \in \delta(k)} (p_{km}^{\mathrm{T}} + jq_{km}^{\mathrm{T}}), \; \textbf{(reported power flow balance in attacked zone)} \tag{1d}$$

$$\forall k \in \mathcal{A}^C \setminus \mathcal{R} : \quad \hat{P}_k^g - \hat{P}_k^d + j(\hat{Q}_k^g - \hat{Q}_k^g) = \sum_{km \in \delta(k)} (p_{km}^{\mathrm{T}} + jq_{km}^{\mathrm{T}}) \; \textbf{(LHS is data, not variables)} \tag{1e}$$

$$\forall k \in \mathcal{R} : \quad P_k^g - \hat{P}_k^d + j(Q_k^g - \hat{Q}_k^g) = \sum_{km \in \delta(k)} (p_{km}^{\mathrm{T}} + jq_{km}^{\mathrm{T}}) \; (\boldsymbol{P_k^g}, \boldsymbol{Q_k^g} \textbf{ are variables}) \tag{1f}$$

$$P_k^g - \hat{P}_k^g = \alpha_k \Delta \qquad \textbf{(AGC response)} \; \boldsymbol{\Delta} \textbf{ is a variable}, \tag{1g}$$

reported data: operational limits on all buses, generators and lines (1h)

all $\boldsymbol{p_{km}^{\mathrm{T}}}, \boldsymbol{q_{km}^{\mathrm{T}}}$ related to $|\boldsymbol{V_k^{\mathrm{T}}}|, |\boldsymbol{V_m^{\mathrm{T}}}|, \boldsymbol{\theta_k^{\mathrm{T}}}, \boldsymbol{\theta_m^{\mathrm{T}}}$ through AC power flow laws (1i)

Undetectable static attack (load modification, no line tripping, abridged!)

$$\textbf{Max} \quad (p_{uv}^{\mathrm{T}})^2 + (q_{uv}^{\mathrm{T}})^2 \quad \text{square norm of flow on uv} \tag{1a}$$

s.t.

$$\forall k \in \mathcal{A}^C \cup \partial\mathcal{A}, \ V_k^{\mathrm{T}} = V_k^{\mathrm{R}} \ \text{(truthful voltages outside attacked zone)} \tag{1b}$$

$$\forall k \in \mathcal{A}, \ -(P_k^{d,\mathrm{R}} + jQ_k^{d,\mathrm{R}}) = \sum_{km \in \delta(k)} (p_{km}^{\mathrm{R}} + jq_{km}^{\mathrm{R}}), \ \text{(true power flow balance in attacked zone)} \tag{1c}$$

$$-(P_k^{d,\mathrm{T}} + jQ_k^{d,\mathrm{T}}) = \sum_{km \in \delta(k)} (p_{km}^{\mathrm{T}} + jq_{km}^{\mathrm{T}}), \ \text{(reported power flow balance in attacked zone)} \tag{1d}$$

$$\forall k \in \mathcal{A}^C \backslash \mathcal{R}: \quad \hat{P}_k^g - \hat{P}_k^d + j(\hat{Q}_k^g - \hat{Q}_k^g) = \sum_{km \in \delta(k)} (p_{km}^{\mathrm{T}} + jq_{km}^{\mathrm{T}}) \ \text{(LHS is data, not variables)} \tag{1e}$$

$$\forall k \in \mathcal{R}: \quad P_k^g - \hat{P}_k^d + j(Q_k^g - \hat{Q}_k^g) = \sum_{km \in \delta(k)} (p_{km}^{\mathrm{T}} + jq_{km}^{\mathrm{T}}) \ (\boldsymbol{P_k^g}, \boldsymbol{Q_k^g} \text{ are variables)} \tag{1f}$$

$$P_k^g - \hat{P}_k^g = \alpha_k \Delta \quad \text{(AGC response)} \ \boldsymbol{\Delta} \text{ is a variable}, \tag{1g}$$

reported data: operational limits on all buses, generators and lines (1h)

all $\boldsymbol{p_{km}^{\mathrm{T}}}, \boldsymbol{q_{km}^{\mathrm{T}}}$ related to $|\boldsymbol{V_k^{\mathrm{T}}}|, |\boldsymbol{V_m^{\mathrm{T}}}|, \boldsymbol{\theta_k^{\mathrm{T}}}, \boldsymbol{\theta_m^{\mathrm{T}}}$ through AC power flow laws (1i)

AC OPF-like problem, local-solvable in seconds

A large-scale example: `case2746wp`

(2746 buses, 3514 lines, 520 generators, 25GW total load)

A large-scale example: `case2746wp`

(2746 buses, 3514 lines, 520 generators, 25GW total load)

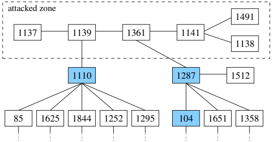

Undetectable attack with strong overloads on branches:

(1361, 1141):
$\|$reported flow$\|$ = 109, $\|$true flow$\|$ = **229**, limit = 114

(1138, 1141):
$\|$reported flow$\|$ = 98, $\|$true flow$\|$ = **209**, limit = 114

Net load change: **135 MW (< 0.5%)** of total load

Non-static attack: follow-up

A blind spot in prior work?

Non-static attack: follow-up

A blind spot in prior work?

"Noisy-data"

Following the attack, for any bus $\in \mathcal{A} - \partial\mathcal{A}$ the attacker reports (at each time t) a complex voltage value

$$\tilde{V}_k(t) \;=\; V_k^{\mathrm{R}} + \boldsymbol{\nu_k(t)}$$

Here, $\boldsymbol{\nu_k(t)}$ is *random*, with

$$\mathrm{E}(\boldsymbol{\nu_k(t)}) = 0,$$

Non-static attack: follow-up

A blind spot in prior work?

"Noisy-data"

Following the attack, for any bus $\in \mathcal{A} - \partial\mathcal{A}$ the attacker reports (at each time t) a complex voltage value

$$\tilde{V}_k(t) \;=\; V_k^{\mathrm{R}} + \boldsymbol{\nu_k(t)}$$

Here, $\boldsymbol{\nu_k(t)}$ is *random*, with

$$\mathrm{E}(\boldsymbol{\nu_k(t)}) = 0,$$

(consistent with zero expected load change)

Non-static attack: follow-up

A blind spot in prior work?

"Noisy-data"

Following the attack, for any bus $\in \mathcal{A} - \partial\mathcal{A}$ the attacker reports (at each time t) a complex voltage value

$$\tilde{V}_k(t) \;=\; V_k^{\mathrm{R}} + \boldsymbol{\nu_k(t)}$$

Here, $\boldsymbol{\nu_k(t)}$ is *random*, with

$$\mathrm{E}(\boldsymbol{\nu_k(t)}) = 0,$$

(consistent with zero expected load change)

Non-static attack: follow-up

A blind spot in prior work?

"Noisy-data"

Following the attack, for any bus $\in \boldsymbol{\mathcal{A} - \partial\mathcal{A}}$ the attacker reports (at each time t) a complex voltage value

$$\tilde{V}_k(t) \;=\; V_k^{\mathrm{R}} + \boldsymbol{\nu_k(t)}$$

Here, $\boldsymbol{\nu_k(t)}$ is *random*, with

$$\mathrm{E}(\boldsymbol{\nu_k(t)}) = 0,$$

(consistent with zero expected load change)

and?

Non-static attack: follow-up

A blind spot in prior work?

"Noisy-data"

Following the attack, for any bus $\in \mathcal{A} - \partial\mathcal{A}$ the attacker reports (at each time t) a complex voltage value

$$\tilde{V}_k(t) \ = \ V_k^{\mathrm{R}} + \boldsymbol{\nu_k(t)}$$

Here, $\boldsymbol{\nu_k(t)}$ is *random*, with

$$\mathrm{E}(\boldsymbol{\nu_k(t)}) = 0,$$

(consistent with zero expected load change)

and? what else?

Non-static attack: follow-up

A blind spot in prior work?

"Noisy-data"

Following the attack, for any bus $\in \mathcal{A} - \partial\mathcal{A}$ the attacker reports (at each time t) a complex voltage value

$$\tilde{V}_k(t) \;=\; V_k^{\mathrm{R}} + \boldsymbol{\nu_k(t)}$$

Here, $\boldsymbol{\nu_k(t)}$ is *random*, with

$$\mathrm{E}(\boldsymbol{\nu_k(t)}) = 0,$$

(consistent with zero expected load change)

and? what else?

Defense, 0

- Defender is likely to know that “something” happened (and quickly). But sensor data is noisy and “something” may be inconsequential

Defense, 0

- Defender is likely to know that “something” happened (and quickly). But sensor data is noisy and “something” may be inconsequential
- We want a defensive action that is nearly implementable in terms of today’s grid operation

Defense, 0

- Defender is likely to know that "something" happened (and quickly). But sensor data is noisy and "something" may be inconsequential
- We want a defensive action that is nearly implementable in terms of today's grid operation
- Should not lead to false positives

Defense, 0

- Defender is likely to know that "something" happened (and quickly). But sensor data is noisy and "something" may be inconsequential
- We want a defensive action that is nearly implementable in terms of today's grid operation
- Should not lead to false positives
- **Solution:** change the power quantities in a way that the attacker cannot anticipate, and identify inconsistent signals. How?

Defense, 0

- Defender is likely to know that "something" happened (and quickly). But sensor data is noisy and "something" may be inconsequential
- We want a defensive action that is nearly implementable in terms of today's grid operation
- Should not lead to false positives
- **Solution:** change the power quantities in a way that the attacker cannot anticipate, and identify inconsistent signals. How?
- A solution: change generator output by a **random** injection that yields a **valid** power flow solution ("AGC-lite" plus redispatch)

Defense, 0' (optimization problem)

Following attack, and in suspicion of an attack

- Defender only has access to **reported** data, which is accurate in the non-attacked zone. But the defender **does not** know the attacked zone.

Defense, 0' (optimization problem)

Following attack, and in suspicion of an attack

- Defender only has access to **reported** data, which is accurate in the non-attacked zone. But the defender **does not** know the attacked zone.
- **(repeatedly)** Defender chooses a random subset of the AGC-responding generators, and

Defense, 0' (optimization problem)

Following attack, and in suspicion of an attack

- Defender only has access to **reported** data, which is accurate in the non-attacked zone. But the defender **does not** know the attacked zone.
- **(repeatedly)** Defender chooses a random subset of the AGC-responding generators, and
- Defender computes a random power flow solution where the chosen generators are allowed to change (up or down) their output, within limits. Other generators can change output by small amounts, within limits. The power flow solution must satisfy e.g. **voltage constraints**.

Defense, 0' (optimization problem)

Following attack, and in suspicion of an attack

- Defender only has access to **reported** data, which is accurate in the non-attacked zone. But the defender **does not** know the attacked zone.
- **(repeatedly)** Defender chooses a random subset of the AGC-responding generators, and
- Defender computes a random power flow solution where the chosen generators are allowed to change (up or down) their output, within limits. Other generators can change output by small amounts, within limits. The power flow solution must satisfy e.g. **voltage constraints**.
- Defender seeks to make the changes in generation large subject to above constraints.

Defense, 0' (optimization problem)

Following attack, and in suspicion of an attack

- Defender only has access to **reported** data, which is accurate in the non-attacked zone. But the defender **does not** know the attacked zone.
- **(repeatedly)** Defender chooses a random subset of the AGC-responding generators, and
- Defender computes a random power flow solution where the chosen generators are allowed to change (up or down) their output, within limits. Other generators can change output by small amounts, within limits. The power flow solution must satisfy e.g. **voltage constraints**.
- Defender seeks to make the changes in generation large subject to above constraints. **ACOPF**-like problem, solvable in seconds

Defense, 1

But attacker cannot anticipate this random action, even if the defense is known.

Defense, 1

But attacker cannot anticipate this random action, even if the defense is known. **Therefore:**

Defense, 1

But attacker cannot anticipate this random action, even if the defense is known. **Therefore:**
(under noisy-data attack)

- Reported currents, and implied power flows, will have **near-constant** values within attacked zone
- But outside of attacked zone, with high-probability (?) most lines will see significant changes in current and power flows

Defense, 1

But attacker cannot anticipate this random action, even if the defense is known. **Therefore:**
(under noisy-data attack)

- Reported currents, and implied power flows, will have **near-constant** values within attacked zone
- But outside of attacked zone, with high-probability (?) most lines will see significant changes in current and power flows

Above example (`case2746wp`) has over **3500** lines, but in a few iterations we reduce the number of suspicious lines to $< \mathbf{100}$.

Defense, 1

But attacker cannot anticipate this random action, even if the defense is known. **Therefore:**
(under noisy-data attack)

- Reported currents, and implied power flows, will have **near-constant** values within attacked zone
- But outside of attacked zone, with high-probability (?) most lines will see significant changes in current and power flows

Above example (`case2746wp`) has over **3500** lines, but in a few iterations we reduce the number of suspicious lines to $< \mathbf{100}$.

Defense, 1

But attacker cannot anticipate this random action, even if the defense is known. **Therefore:**
(under noisy-data attack)

- Reported currents, and implied power flows, will have **near-constant** values within attacked zone
- But outside of attacked zone, with high-probability (?) most lines will see significant changes in current and power flows

Above example (`case2746wp`) has over **3500** lines, but in a few iterations we reduce the number of suspicious lines to $< \mathbf{100}$.

Good, but not good enough

Defense, 2: $\begin{pmatrix} I_{km} \\ I_{mk} \end{pmatrix} = Y_{km} \begin{pmatrix} V_k \\ V_m \end{pmatrix}$

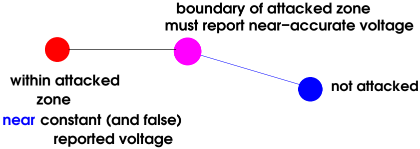

Defense, 2: $\begin{pmatrix} I_{km} \\ I_{mk} \end{pmatrix} = Y_{km} \begin{pmatrix} V_k \\ V_m \end{pmatrix}$

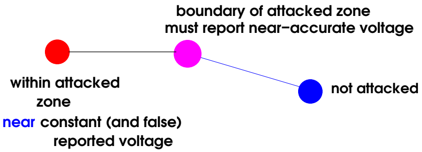

On a line going from boundary to interior of attacked zone

Defense, 2: $\begin{pmatrix} I_{km} \\ I_{mk} \end{pmatrix} = Y_{km} \begin{pmatrix} V_k \\ V_m \end{pmatrix}$

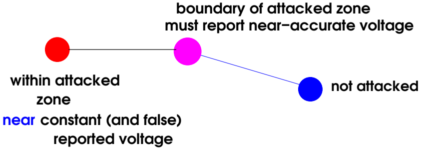

On a line going from boundary to interior of attacked zone **reported** current will be wrong

Defense, 2: $\begin{pmatrix} I_{km} \\ I_{mk} \end{pmatrix} = Y_{km} \begin{pmatrix} V_k \\ V_m \end{pmatrix}$

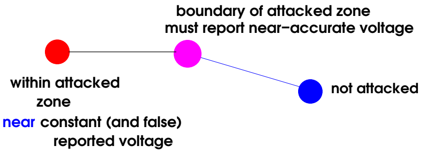

On a line going from boundary to interior of attacked zone **reported** current will be wrong

because voltage at boundary bus is changing with our defense

Defense, 2: $\begin{pmatrix} I_{km} \\ I_{mk} \end{pmatrix} = Y_{km} \begin{pmatrix} V_k \\ V_m \end{pmatrix}$

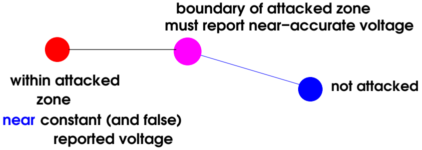

On a line going from boundary to interior of attacked zone **reported** current will be wrong

because voltage at boundary bus is changing with our defense but voltage at interior bus is changing by very small amounts

Defense, 2: $\begin{pmatrix} I_{km} \\ I_{mk} \end{pmatrix} = Y_{km} \begin{pmatrix} V_k \\ V_m \end{pmatrix}$

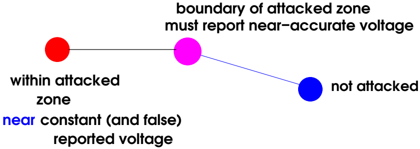

On a line going from boundary to interior of attacked zone
reported current will be wrong

because voltage at boundary bus is changing with our defense
but voltage at interior bus is changing by very small amounts

In above example, **one** iteration identifies all boundary lines with **no** false positives

	Experiment 1	Experiment 2
$\sum_{j\in\mathcal{G}} \delta_j^+$	289.01	964.77
$\sum_{j\in\mathcal{G}} \delta_j^-$	174.47	256.04

Branch $(k = 1137, m = 1139)$

1137 **inside attack**, 1139 **on boundary**

$\lvert V^{\mathrm{R}}_{1137}(0)\rvert\angle\theta^{\mathrm{R}}_{1137}(0)$	$1.0919\angle-6.993^\circ$	$1.0919\angle-6.993^\circ$
$\lvert V^{\mathrm{R}}_{1139}(0)\rvert\angle\theta^{\mathrm{R}}_{1139}(0)$	$1.0919\angle-6.991^\circ$	$1.0919\angle-6.991^\circ$
$\lvert V^{\mathrm{R}}_{1139}(t)\rvert\angle\theta^{\mathrm{R}}_{1139}(t)$	$\mathbf{1.0105\angle-7.882^\circ}$	$\mathbf{1.0187\angle-7.936^\circ}$
$I^{\mathrm{R}}_{1137,1139}(0)$	$-0.0275+0.0281j$	$-0.0275+0.0281j$
$Y_{1137,1139}\begin{pmatrix}V^{\mathrm{R}}_{1137}(0)\\ V^{\mathrm{R}}_{1139}(t)\end{pmatrix}$	$\mathbf{20.967-55.978j}$	$\mathbf{21.435-49.918j}$

Non-static attack: follow-up

“Noisy-data” attack

Following the attack, for any bus $\in \boldsymbol{\mathcal{A} - \partial\mathcal{A}}$ the attacker reports (at each time t) a complex voltage value

$$\tilde{V}_k(t) \; = \; V_k^{\mathrm{R}} + \boldsymbol{\nu_k(t)}$$

Here, $\boldsymbol{\nu_k(t)}$ is *random*, with

$$\mathrm{E}(\boldsymbol{\nu_k(t)}) = 0,$$

Non-static attack: follow-up

"Noisy-data" attack

Following the attack, for any bus $\in \mathcal{A} - \partial\mathcal{A}$ the attacker reports (at each time t) a complex voltage value

$$\tilde{V}_k(t) \; = \; V_k^{\mathrm{R}} + \boldsymbol{\nu_k(t)}$$

Here, $\boldsymbol{\nu_k(t)}$ is *random*, with

$$\mathrm{E}(\boldsymbol{\nu_k(t)}) = 0,$$

(consistent with zero expected load change)

Non-static attack: follow-up

"Noisy-data" attack

Following the attack, for any bus $\in \mathcal{A} - \partial\mathcal{A}$ the attacker reports (at each time t) a complex voltage value

$$\tilde{V}_k(t) \;=\; V_k^{\mathrm{R}} + \boldsymbol{\nu_k(t)}$$

Here, $\boldsymbol{\nu_k(t)}$ is *random*, with

$$\mathrm{E}(\boldsymbol{\nu_k(t)}) = 0,$$

(consistent with zero expected load change)
and?

Non-static attack: follow-up

“Noisy-data” attack

Following the attack, for any bus $\in \mathcal{A} - \partial\mathcal{A}$ the attacker reports (at each time t) a complex voltage value

$$\tilde{V}_k(t) \; = \; V_k^{\mathrm{R}} + \boldsymbol{\nu_k(t)}$$

Here, $\boldsymbol{\nu_k(t)}$ is *random*, with

$$\mathrm{E}(\boldsymbol{\nu_k(t)}) = 0,$$

(consistent with zero expected load change)
and? what else?

Non-static attack: follow-up

"Noisy-data" attack

Following the attack, for any bus $\in \mathcal{A} - \partial\mathcal{A}$ the attacker reports (at each time t) a complex voltage value

$$\tilde{V}_k(t) \; = \; V_k^{\mathrm{R}} + \boldsymbol{\nu_k(t)}$$

Here, $\boldsymbol{\nu_k(t)}$ is *random*, with

$$\mathrm{E}(\boldsymbol{\nu_k(t)}) = 0,$$

(consistent with zero expected load change)
and? what else?

$\rightarrow$ *stochastics* of $\boldsymbol{\nu_k(t)}$ should "make sense"

PMU fun

We have data from an industrial partner:

- 240 PMUs
- 2 years of reported data
- 28 TB
- Soon, 500 PMUs and higher detail

PMU fun

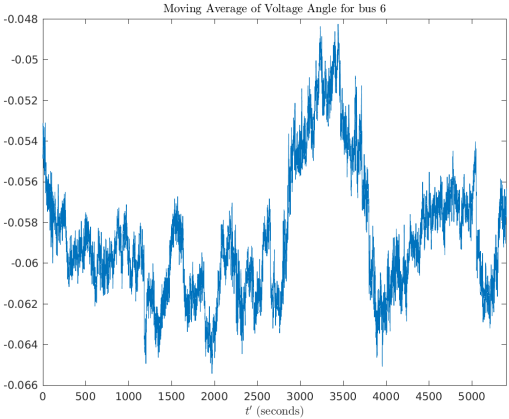

Moving Average of Voltage Angle for bus 6
-0.048
-0.05
-0.052
-0.054
-0.056
-0.058
-0.06
-0.062
-0.064
-0.066
0
500
1000
1500
2000
2500
3000
3500
4000
4500
5000
t' (seconds)

More PMU fun: a voltage phase angle

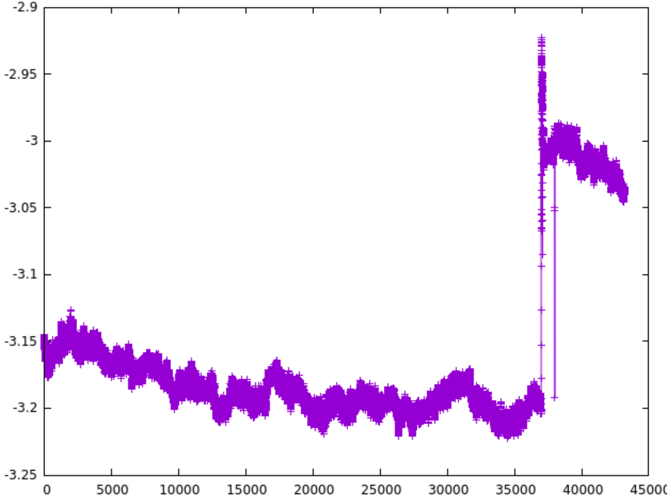

More PMU fun: 3 voltage angles)

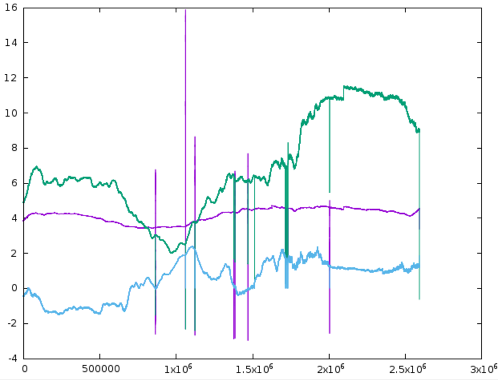

More PMU fun: difference between two voltage angles (10 seconds)

More PMU fun: frequency at two different buses

Noise is not just noise

From real time series, voltage angle deviation histogram

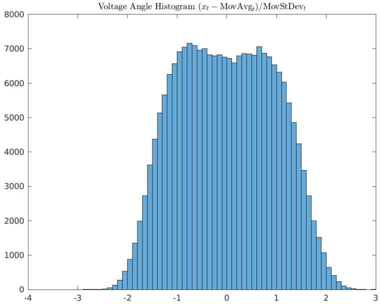

Kolmogorov-Smirnoff gaussianity test strongly rejected, always

Noise is not just noise

From real time series, voltage magnitude deviations

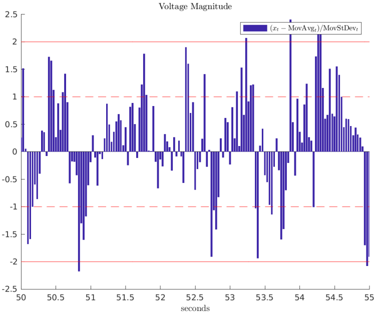

Strong and nontrivial correlation structure

Covariances matrices of PMU data have **low rank!!**

Covariances matrices of PMU data have **low rank!!**

Example: 50 PMUs, Voltage Angle, one minute

	Scaled Eigenvalue
1	1.000
2	0.078
3	0.012
4	0.009
5	0.007
6	0.004
7	0.003
8	0.002
9	0.001
10	0.001

Covariances matrices of PMU data have **low rank!!**

Covariances matrices of PMU data have **low rank!!**

Example: 100 PMUs, voltage magnitude, five minutes

	Scaled Eigenvalue
1	1.000
2	0.618
3	0.061
4	0.023
5	0.017
6	0.010
7	0.008
8	0.004
9	0.004
10	0.002

Non-static attack: follow-up

"Noisy-data" attack

Following the attack, for any bus $\boldsymbol{k} \in \mathcal{A} - \partial\mathcal{A}$ the attacker reports (at each time t) a complex voltage value

$$\tilde{V}_k^{\mathrm{R}}(t) \; = \; V_k^{\mathrm{R}} + \boldsymbol{\nu_k(t)}$$

Here, $\boldsymbol{\nu_k(t)}$ is *random*, with

$$\mathrm{E}(\boldsymbol{\nu_k(t)}) = 0,$$

Non-static attack: follow-up

"Noisy-data" attack

Following the attack, for any bus $\boldsymbol{k} \in \mathcal{A} - \partial\mathcal{A}$ the attacker reports (at each time t) a complex voltage value

$$\tilde{V}_k^{\mathrm{R}}(t) \;=\; V_k^{\mathrm{R}} + \boldsymbol{\nu_k(t)}$$

Here, $\boldsymbol{\nu_k(t)}$ is *random*, with

$$\mathrm{E}(\boldsymbol{\nu_k(t)}) = 0,$$

(consistent with zero expected load change)

Non-static attack: follow-up

"Noisy-data" attack

Following the attack, for any bus $\boldsymbol{k} \in \boldsymbol{\mathcal{A}} - \boldsymbol{\partial\mathcal{A}}$ the attacker reports (at each time t) a complex voltage value

$$\tilde{V}_k^{\mathrm{R}}(t) \ = \ V_k^{\mathrm{R}} + \boldsymbol{\nu_k}(\boldsymbol{t})$$

Here, $\boldsymbol{\nu_k}(\boldsymbol{t})$ is *random*, with

$$\mathrm{E}(\boldsymbol{\nu_k}(\boldsymbol{t})) = 0,$$

(consistent with zero expected load change)
and?

Non-static attack: follow-up

"Noisy-data" attack

Following the attack, for any bus $\boldsymbol{k} \in \mathcal{A} - \partial\mathcal{A}$ the attacker reports (at each time t) a complex voltage value

$$\tilde{V}_k^{\mathrm{R}}(t) \; = \; V_k^{\mathrm{R}} + \boldsymbol{\nu_k}(\boldsymbol{t})$$

Here, $\boldsymbol{\nu_k}(\boldsymbol{t})$ is *random*, with

$$\mathrm{E}(\boldsymbol{\nu_k}(\boldsymbol{t})) = 0,$$

(consistent with zero expected load change)
and? what else?

Non-static attack: follow-up

"Noisy-data" attack

Following the attack, for any bus $\boldsymbol{k} \in \mathcal{A} - \partial\mathcal{A}$ the attacker reports (at each time t) a complex voltage value

$$\tilde{V}_k^{\mathrm{R}}(t) \;=\; V_k^{\mathrm{R}} + \boldsymbol{\nu_k}(\boldsymbol{t})$$

Here, $\boldsymbol{\nu_k}(\boldsymbol{t})$ is *random*, with

$$\mathrm{E}(\boldsymbol{\nu_k}(\boldsymbol{t})) = 0,$$

(consistent with zero expected load change)
and? what else?

$\rightarrow$ *covariance* of $\boldsymbol{\nu}(\boldsymbol{t})$ should be **make sense**

Learning variances

Theorem. (Co)variance of time series can be learned

- In real time
- In streaming fashion
- Under evolving stochasticity

Shukla, Yun ~~and a fool from Columbia~~ :
Non-Stationary Streaming PCA, Proc. 2017 NIPS Time Series Workshop.

Covariance defense

- Under **whatever** assumptions, the attacker will produce a time series for e.g. phase angles.

Covariance defense

- Under **whatever** assumptions, the attacker will produce a time series for e.g. phase angles.
- Assume covariance of phase angles is learned by the defender

Covariance defense

- Under **whatever** assumptions, the attacker will produce a time series for e.g. phase angles.
- Assume covariance of phase angles is learned by the defender
- (Assume of low rank)

Covariance defense

- Under **whatever** assumptions, the attacker will produce a time series for e.g. phase angles.
- Assume covariance of phase angles is learned by the defender
- (Assume of low rank)
- Defender chooses **random generator injections so as to significantly change covariance of phase angles**

Covariance defense

- Under **whatever** assumptions, the attacker will produce a time series for e.g. phase angles.
- Assume covariance of phase angles is learned by the defender
- (Assume of low rank)
- Defender chooses **random generator injections so as to significantly change covariance of phase angles**
- Attacker is caught with pants down

Covariance defense (technical, abridged)

- Let Ω = covariance of **observed** voltage phase angles

Covariance defense (technical, abridged)

- Let $\color{red}{\boldsymbol{\Omega}}$ = covariance of **observed** voltage phase angles

- Let $\boldsymbol{w_1}, \boldsymbol{w_2}, \ldots, \boldsymbol{w_r}$ = eigenvectors with ~~positive~~ large enough eigenvalues.

Covariance defense (technical, abridged)

- Let Ω = covariance of **observed** voltage phase angles
- Let $\boldsymbol{w_1}, \boldsymbol{w_2}, \ldots, \boldsymbol{w_r}$ = eigenvectors with ~~positive~~ large enough eigenvalues. $\boldsymbol{r} \ll \boldsymbol{n}$ (number of buses)

Covariance defense (technical, abridged)

- Let $\boldsymbol{\Omega}$ = covariance of **observed** voltage phase angles

- Let $\boldsymbol{w_1}, \boldsymbol{w_2}, \ldots, \boldsymbol{w_r}$ = eigenvectors with ~~positive~~ large enough eigenvalues. $\boldsymbol{r} \ll \boldsymbol{n}$ (number of buses)

- Defender chooses vector $\boldsymbol{v} \in \mathbb{R}^{\boldsymbol{n}}$ with:

  $\boldsymbol{w_i^T v = 0}$ for $1 \leq i \leq r$ (plus **other** conditions)

Covariance defense (technical, abridged)

- Let $\boldsymbol{\Omega}$ = covariance of **observed** voltage phase angles

- Let $\boldsymbol{w_1}, \boldsymbol{w_2}, \ldots, \boldsymbol{w_r}$ = eigenvectors with ~~positive~~ large enough eigenvalues. $\boldsymbol{r} \ll \boldsymbol{n}$ (number of buses)

- Defender chooses vector $\boldsymbol{v} \in \mathbb{R}^{\boldsymbol{n}}$ with:

  $\boldsymbol{w_i^T v = 0}$ for $1 \leq i \leq r$ (plus **other** conditions)

- **Theorem:** there is a random set of power injections (by generators) that results in covariance of phase angles

  $$\approx \; \boldsymbol{\Omega} + \boldsymbol{\lambda v v^T}$$

Covariance defense (technical, abridged)

- Let $\boldsymbol{\Omega}$ = covariance of **observed** voltage phase angles
- Let $\boldsymbol{w_1}, \boldsymbol{w_2}, \ldots, \boldsymbol{w_r}$ = eigenvectors with ~~positive~~ large enough eigenvalues. $\boldsymbol{r} \ll \boldsymbol{n}$ (number of buses)
- Defender chooses vector $\boldsymbol{v} \in \mathbb{R}^{\boldsymbol{n}}$ with:

  $\boldsymbol{w_i^T v = 0}$ for $1 \leq i \leq r$ (plus **other** conditions)

- **Theorem:** there is a random set of power injections (by generators) that results in covariance of phase angles

$$\approx \ \boldsymbol{\Omega} + \boldsymbol{\lambda v v^T} \text{ where } \boldsymbol{\lambda} > \boldsymbol{0}$$

Covariance defense (technical, abridged)

- Let $\boldsymbol{\Omega}$ = covariance of **observed** voltage phase angles
- Let $\boldsymbol{w_1}, \boldsymbol{w_2}, \ldots, \boldsymbol{w_r}$ = eigenvectors with ~~positive~~ large enough eigenvalues. $\boldsymbol{r} \ll \boldsymbol{n}$ (number of buses)
- Defender chooses vector $\boldsymbol{v} \in \mathbb{R}^{\boldsymbol{n}}$ with:

  $\boldsymbol{w_i}^{\boldsymbol{T}} \boldsymbol{v} = \boldsymbol{0}$ for $1 \leq i \leq r$ (plus **other** conditions)
- **Theorem:** there is a random set of power injections (by generators) that results in covariance of phase angles

  $$\approx \; \boldsymbol{\Omega} + \boldsymbol{\lambda} \boldsymbol{v}\boldsymbol{v}^{\boldsymbol{T}} \text{ where } \boldsymbol{\lambda} > \boldsymbol{0}$$
- On `case2746wp`, $\approx$ 10 vectors $\boldsymbol{v}$ cover all buses. (Dense null space vector computation: LP heuristic)

Covariance defense (technical, less abridged)

1. Let $\boldsymbol{\Omega}$ = covariance of **observed** voltage phase angles

2. Let $\boldsymbol{w_1}, \boldsymbol{w_2}, \ldots, \boldsymbol{w_r}$ = eigenvectors with ~~positive~~ large enough eigenvalues.

Covariance defense (technical, less abridged)

1. Let $\boldsymbol{\Omega}$ = covariance of **observed** voltage phase angles

2. Let $\boldsymbol{w_1}, \boldsymbol{w_2}, \ldots, \boldsymbol{w_r}$ = eigenvectors with ~~positive~~ large enough eigenvalues. $\boldsymbol{r} \ll \boldsymbol{n}$ (number of buses)

3. Defender chooses vector $\boldsymbol{v} \in \mathbb{R}^{\boldsymbol{n}}$ with:

   $\boldsymbol{w_i^T v = 0}$ for $1 \leq i \leq r$ and $\boldsymbol{[Bv]_i = 0}$ for all non-generator $\boldsymbol{i}$

4. **Theorem:** there is a random set of power injections (by generators) that results in covariance of phase angles

   $$\approx \boldsymbol{\Omega} + \boldsymbol{\lambda v v^T}$$

Covariance defense (technical, less abridged)

1. Let $\boldsymbol{\Omega}$ = covariance of **observed** voltage phase angles

2. Let $\boldsymbol{w_1}, \boldsymbol{w_2}, \ldots, \boldsymbol{w_r}$ = eigenvectors with ~~positive~~ large enough eigenvalues. $\boldsymbol{r} \ll \boldsymbol{n}$ (number of buses)

3. Defender chooses vector $\boldsymbol{v} \in \mathbb{R}^{\boldsymbol{n}}$ with:

   $\boldsymbol{w_i^T v = 0}$ for $1 \le i \le r$ and $\boldsymbol{[Bv]_i = 0}$ for all non-generator $\boldsymbol{i}$

4. **Theorem:** there is a random set of power injections (by generators) that results in covariance of phase angles

   $$\approx \boldsymbol{\Omega + \lambda v v^T} \text{ where } \boldsymbol{\lambda > 0}$$

5. On `case2746wp`, there is a **single** vector $\boldsymbol{v}$ that covers all buses.

Covariance defense (technical, less abridged)

1. Let $\boldsymbol{\Omega}$ = covariance of **observed** voltage phase angles

2. Let $\boldsymbol{w_1}, \boldsymbol{w_2}, \ldots, \boldsymbol{w_r}$ = eigenvectors with ~~positive~~ large enough eigenvalues. $\boldsymbol{r} \ll \boldsymbol{n}$ (number of buses)

3. Defender chooses vector $\boldsymbol{v} \in \mathbb{R}^{\boldsymbol{n}}$ with:

   $\boldsymbol{w_i^T v = 0}$ for $1 \leq i \leq r$ and $\boldsymbol{[Bv]_i = 0}$ for all non-generator $\boldsymbol{i}$

4. **Theorem:** there is a random set of power injections (by generators) that results in covariance of phase angles

   $$\approx \ \boldsymbol{\Omega} + \boldsymbol{\lambda v v^T} \text{ where } \boldsymbol{\lambda > 0}$$

5. On `case2746wp`, there is a **single** vector $\boldsymbol{v}$ that covers all buses.

   Theorem: if v^1, v^2 $\in$ subspace S, then $\exists \, \infty$ many $v \in S$ with

   $$\text{support}(v) = \text{support}(v^1) \cup \text{support}(v^2)$$

Summary

- Very high-fidelity grid attacks appear easily computable.
- Defensive idea **1**: use network resources to change power flow physics in unpredictable ways
- Defensive idea **2**: change covariance structure in a way that cannot be instantaneously learned

Summary

- Very high-fidelity grid attacks appear easily computable.
- Defensive idea **1**: use network resources to change power flow physics in unpredictable ways
- Defensive idea **2**: change covariance structure in a way that cannot be instantaneously learned
- Adversarial learning of moments under streaming data is a nice problem!

Wed.Aug.15.113946.2018@blacknwhite