Security and Statistics on Power Grids

Mauro Cesar Escobar Santoro

Submitted in partial fulfillment of the
requirements for the degree
of Doctor of Philosophy
in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2019

©2019
Mauro Cesar Escobar Santoro
All Rights Reserved

ABSTRACT

Security and Statistics on Power Grids

Mauro Cesar Escobar Santoro

Improving the functioning and the safety of the electrical grids is a topic of great concern, given its magnitude and importance in today's world. In this thesis, we focus in these two subjects.

In the first part, we study undetectable cyber-physical attacks on power grids, which are attacks that involve physical disruptions, including tripping lines and load modifications, and sensor output alterations. We propose a sophisticated attack model described under the full Alternating Current (AC) power flow equations and show its feasibility on large grids from a test cases library. As counter-measures, we propose different defensive strategies that the network's controller can apply under a suspected cyber attack. These are random, simple and fast procedures that change the voltages across the network and aim to unmask the current status of the system, assuming that the attacker cannot react against their randomness.

Secondly, with access to data collected through Phasor Measurement Units (PMUs) by a power utility in the United States, we perform statistical analyses on the frequency and voltage time series that have been recorded at a rate of 30 Hz. We focus on intervals of time where the sampled data shows to be in steady-state conditions and, with the use of appropriate signal processing filters, we are able to extract hidden anomalies such as spatio-temporal correlations between sensors and harmonic distortions.

Table of Contents

List of Figures

List of Tables

Acknowledgments

I would like to thank my advisor, Prof. Daniel Bienstock, for being a great and dedicated mentor and from whom I have learned a lot during my time at Columbia University. Thanks to him I started to explore the world of optimization applied to the power grid, without previous knowledge of electricity. I feel very fortunate of receiving his devotion to instructing me through the research that we have done.

I would also like to thank the wisdom, expertise and contribution that I have received from faculty members of the department, specially from Profs. Garud Iyengar, Ali Hirsa, Yuri Faenza, and Donald Goldfarb. I have truly enjoyed their classes and I appreciate every comment and feedback that they have given me. I believe that, together with Prof. Daniel Bienstock, they are a central piece in the optimization community of the department.

I want to thank Prof. Michael Chertkov that hosted me for three weeks in Los Alamos National Lab. Thanks to him, my time there was very productive and I could get a new perspective of the analyses that we were doing. In addition, I am grateful for the collaboration held with Dr. Jonatan Ostrometzky and Prof. Gil Zussman, from the Electrical Engineering department. The Fourier analysis that we did are based on their expertise and feedback.

IEOR has been a wonderful place to be part of, I will never forget the people that I met in the department. Thanks to them, which many have become truly friends, the years that I spent in New York City has been pleasant and enjoyable. Special mention to Gonzalo, Camilo, Liz, Enrique, Julien, and Nouri.

I want to thank Eduard for the company through most of this process, I receive

constant motivation and energy from admiring what he does, I am very grateful for what we have. Lastly, to my family for being always supportive and encouraging, I constantly feel them close and present even though they are far away. They are a fundamental pillar of who I am and the path that I have followed.

Chapter 1

Introduction

The electrical grid is one the largest and most complex infrastructures built by mankind. It has a great importance in our lives, moreover, we have become daily dependent on it: most jobs make use of electricity, as well as modern systems of telecommunication, means of transportation, and important institutions —such as governments or hospitals. In summary, the world's economies rely on a solid system to supply power. Because of this reason, it is crucial to maintain a stable and safe system that provides energy to its customers in a reliable manner.

The existing power grid in the United States nowadays presents several challenges. It has exponentially grown across the last century without planning its future development and, consequently, faces weaknesses when it has to satisfy the increasing demand seen nowadays. As detailed in Gretchen Bakke's book [10] from 2016, more than 70 percent of the transmission lines and transformers in the United States are twenty-five years old and the average age of the power plants is thirty-four. As a result, the system has large inefficiencies that have caused an increasing number of power outages across the last decades: 15 in 2001, 78 in 2007, 307 in 2011. Additionally, the unpredictability of renewable energy sources has contributed to the instability of the grid —where production has to meet demand at all times. The current grid is not prepared to deal with this uncertainty.

These current challenges together with the complexity of the physics that supports the generation and transmission of electricity throughout a system has motivated decades of research on different topics related with the power network. The equations behind the model of the power flows are not simple —they are, in its full and most realistic form, non-linear and non-convex— and, therefore, difficult to understand and compute with. Nevertheless, decades of investigation have lead us to have a trustworthy system that works well most of the time.

The main goal of this thesis is to contribute towards a better understanding of the power grid's functioning and improving its safety. For that purpose we first study cyber-physical attacks on the electrical grid, meaning that there is an adversary that has the potential ability to cause physical damage or disruption in the system —such as, a tripping a line, modifying demand in certain points, or changing the functioning of a transformer— coordinated with a cyber attack. A cyber attack corresponds to a hack that blocks or alters the data that the grid's controller obtains from sensors distributed at several locations in the network. These sensors provide information used to estimate the status and the state variables of the system and, this system is used by the controller to make control decisions. The purpose of the hack perpetrated by the adversary, is to hide from the controller his or her physical actions (the physical attack) in order that the controller remains unaware of them and, potentially, executes inopportune decisions.

Diverse models of cyber and cyber-physical attacks have been proposed and studied in the last years, most of them using a linear approximation of the power flow equations. In this thesis, we analyze a new model where the adversary can modify demand of a subset of nodes of the grid and at the same time alter the measurements sampled in that zone. We use the full description of the power flow equations and show that the attack can cause important overloads that, if kept hidden from the controller for long enough time, would have unfavorable consequences.

We complement this section by also describing different strategies that the grid's con-

troller (under the suspicion of an attack) can perform in order to unmask the real status of the variables of the system. These strategies are simple generic procedures that can be used in different kind of cyber attacks.

As a second theme of research, we center our attention on the measurements sampled over the grid. With new technologies and the digital age, modernization of equipment and control of the network has improved, which is giving an opportunity to enhance processes and have a better understanding of the network behavior. One particular example is the upgrade of the sensors that measure electro-physical quantities across the grid (voltage, frequency, current): from the old RTUs (remote terminal units) that produce samples every few seconds to the new PMUs (phasor measurement units) that sample as fast as 30 times per second. With this new volume of information being generated, we might be able to get a new insight about the nature of the power flowing through the network.

We have at our disposal historical PMU measurements that cover a period of 15 months obtained from a US Independent System Operator; these data correspond to sampling collected in approximately 200 locations at a rate of 30 measurements per second. With this information we performed statistical analyzes using covariance (correlation) matrices in order to understand several questions regarding the structure of the data: Can we obtain useful information from the spetral decomposition of these matrices? Or is the data temporarily correlated? We aim to answer these question by observing graphical results of the statistical tools over selected periods of measurements.

This document is organized as follows: we begin by describing the basic properties and models of the electrical grid and its power flow equations in Chapter 2, we continue with the cyber-attack study in Chapter 3, and finalize with the statistical analysis of the PMU sampling in Chapter 4.

Chapter 2

Preliminaries

In this section we provide a description of a power system that we will use in the next chapters. The mathematical model that is presented is widely used in the literature.

2.1 Power Grids

A power transmission system can be characterized by a graph, where the set of nodes or *buses* represent the physical locations where the power is generated, consumed or redistributed (see [13, 41]). We denote by $\mathcal{N}$ this set of buses and distinguish the set $\mathcal{G} \subset \mathcal{N}$ of *generators*; buses that do not generate power are called *load* buses. The set $\mathcal{E}$ of edges of the graph, also called *branches*, represent the *transmission lines* of the network. For bus $k \in \mathcal{N}$, we denote by $\delta(k) \subset \mathcal{E}$ the set of branches that are incident with k. For a set $\mathcal{A} \subset \mathcal{N}$ of buses, $\mathcal{A}^C$ denotes the complement of $\mathcal{A}$ with respect to $\mathcal{N}$, that is $\mathcal{N} \backslash \mathcal{A}$, and $\partial\mathcal{A}$ denotes the set of buses in $\mathcal{A}$ that are connected with some bus in $\mathcal{A}^C$, i.e. the *boundary* of $\mathcal{A}$. Notation: when x is a complex number, $\overline{x}$ will denote its complex conjugate; and if X is a complex matrix, X^{H} denotes its Hermitian transpose (i.e. conjugate transpose, $\overline{X^\top}$).

While complex power is generated and consumed throughout the system, transmission lines are in charge of carrying the power across the different locations. In order to

understand the equations that describe a power flow system, we need to introduce three physical concepts present in the model: voltage, current and power.

The amount of power flowing from bus k to bus m connected by the branch km depends on the physical characteristics of the branch and the voltages V_k and V_m, respectively associated with both buses. Voltages at buses are the *state variables* of the system, they are the result of the whole network configuration (physical characteristics, generation, demand). The voltage at bus k is represented as a complex number, in any of the three different forms:

$$V_k = \underbrace{|V_k|e^{j\theta_k}}_{\text{exponential}} = \underbrace{|V_k|\angle\theta_k}_{\text{polar}} = \underbrace{|V_k|\cos(\theta_k) + j|V_k|\sin(\theta_k)}_{\text{rectangular}}, \tag{2.1}$$

where $|V_k|$ is the magnitude of the voltage and θ_k is its (phase) angle. Voltage, also called *electric potential difference*, makes sense when it is compared between two points or with respect to a reference point, this difference is measured in *volts* or *[V]*. One volt (as defined in [48]) "is the potential difference between two points of a conducting wire carrying a constant current of one ampere, when the power dissipated between these points is equal to one watt."

The *electric current* —or simply *current*— is the flow of electric charge, in our case electrons, moving in a wire. The unit for measurement is the *ampere* or *[A]* defined as "constant current which, if maintained in two straight parallel conductors of infinite length, of negligible circular cross-section, and placed one metre apart in vacuum, would produce between these conductors a force equal to $2 \cdot 10^{-7}$ *[mks]* unit of force (newton) per metre of length" [48]. We denote by I_{km} the complex current that flows from bus k to bus m through branch km.

Finally, the *electric power* is the rate, per unit time, at which the electric energy is transferred through a transmission line. The unit of measurement is the *watt* or *[W]* that "is the power which in one second gives rise to energy of one joule" [48]. We denote by $S_{km} = P_{km} + jQ_{km}$ the complex power that flows from bus k to bus m, where the real part P_{km} is called *active power* and the imaginary part Q_{km} is known as *reactive power*.

It is usually convenient to work with normalized units for the previously mentioned quantities. For this purpose, we define a *per-unit* normalization as

$$\text{quantity in per unit} \;=\; \frac{\text{actual quantity}}{\text{base value of quantity}}, \tag{2.2}$$

where the *base value* for quantities are picked to satisfy the same of relationship as the actual variables [13].

Before making explicit the power flow equations we need to model the branches of the system.

2.1.1 Transmission Lines

We now describe a model that characterizes the physical attributes of the transmission lines (see [18, 101]), including the presence of transformers and shunt admittance. This is also known as π-model.

Consider a branch between buses k and m, see Figure 2.1. Safety operational functioning indicates that the difference of voltage phase angles $\theta_{km} \doteq \theta_k - \theta_m$ between buses k and m should not larger that a pre-established bounds $\theta_{km}^{\min}$ and $\theta_{km}^{\max}$. In addition, branch km has a series *impedance*

$$z_{km} \;=\; r_{km} + jx_{km} \tag{2.3}$$

—that measures the opposition to the flow of current, where r_{km} is the *resistance* and x_{km} is the *reactance*— and series *admittance*

$$y_{km} \;=\; (z_{km})^{-1} \;=\; g_{km} + jb_{km}, \tag{2.4}$$

where $g_{km} = r_{km}/(r_{km}^2 + x_{km}^2)$ is the *conductance* and $b_{km} = -x_{km}/(r_{km}^2 + x_{km}^2)$ is the *susceptance*. There is also a *shunt* —such as a capacitor or inductor, that creates low-resistance— with shunt admittance

$$y_{km}^{\text{sh}} \;=\; g_{km}^{\text{sh}} + jb_{km}^{\text{sh}} \tag{2.5}$$

and a transformer with ratio

$$N_{km} \;=\; \tau_{km} e^{j\sigma_{km}}, \tag{2.6}$$

where $\tau_{km} > 0$ is its magnitude and σ_{km} is the phase shift angle. The transformer is located on the branch next to bus k, it scales the voltages by a factor of $1/N_{km}$ and currents by a factor of $\overline{N_{km}}$.

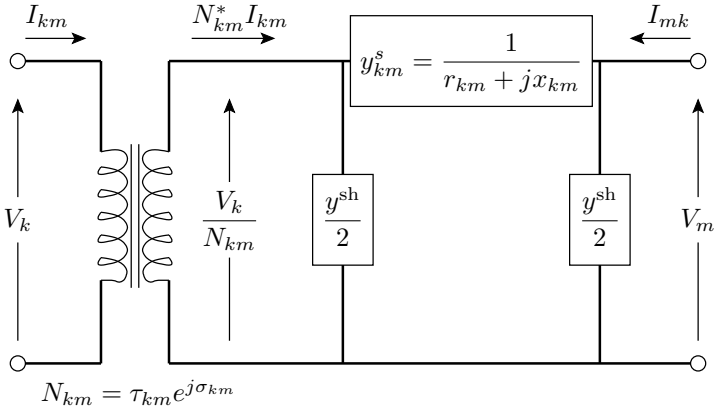

Figure 2.1: Branch Model.

In the rest of this section we will omit the subscript km so as to simplify the equations. With these elements, the *branch admittance matrix* is defined as

$$Y_{km} \;=\; \begin{bmatrix} \mathrm{y}_{kk} & \mathrm{y}_{km} \\ \mathrm{y}_{mk} & \mathrm{y}_{mm} \end{bmatrix} \;=\; \begin{bmatrix} \left(y + \dfrac{y^{\text{sh}}}{2}\right)\dfrac{1}{\tau^2} & -y\dfrac{1}{\tau e^{-j\sigma}} \\ -y\dfrac{1}{\tau e^{j\sigma}} & y + \dfrac{y^{\text{sh}}}{2} \end{bmatrix}. \tag{2.7}$$

Therefore, the complex current injections I_{km} and I_{mk} at buses k and m of the branch, respectively, can be expressed in terms of the voltages V_k and V_m and the admittance matrix Y_{km}:

$$\begin{pmatrix} I_{km} \\ I_{mk} \end{pmatrix} \;=\; Y_{km} \begin{pmatrix} V_k \\ V_m \end{pmatrix} \;=\; \begin{pmatrix} \left(y + \dfrac{y^{\text{sh}}}{2}\right)\dfrac{1}{\tau^2} V_k - y\dfrac{1}{\tau e^{-j\sigma}} V_m \\ -y\dfrac{1}{\tau e^{j\sigma}} V_k + \left(y + \dfrac{y^{\text{sh}}}{2}\right) V_m \end{pmatrix}. \tag{2.8}$$

According to Omh's law, the power injected at a branch is defined as the voltage of

the bus times the complex conjugate of the injected current, in other words:

$$S_{km} = V_k\overline{I_{km}} = \Big(\overline{y} + \frac{\overline{y^{\text{sh}}}}{2}\Big)\frac{1}{\tau^2}|V_k|^2 - \overline{y^s}\frac{1}{\tau e^{j\sigma}}V_k\overline{V_m}, \tag{2.9a}$$

$$S_{mk} = V_m\overline{I_{mk}} = -\overline{y}\frac{1}{\tau e^{-j\sigma}}\overline{V_k}V_m + \Big(\overline{y} + \frac{\overline{y^{\text{sh}}}}{2}\Big)|V_m|^2, \tag{2.9b}$$

from where we can deduce expressions for active and reactive power injected at both ends:

$$P_{km} = \Big(g + \frac{g^{\text{sh}}}{2}\Big)\frac{|V_k|^2}{\tau^2} - \frac{|V_k|}{\tau}|V_m|\big(g\cos(\theta_{km}-\sigma) + b\sin(\theta_{km}-\sigma)\big), \tag{2.10a}$$

$$Q_{km} = -\Big(b + \frac{b^{\text{sh}}}{2}\Big)\frac{|V_k|^2}{\tau^2} + \frac{|V_k|}{\tau}|V_m|\big(b\cos(\theta_{km}-\sigma) - g\sin(\theta_{km}-\sigma)\big), \tag{2.10b}$$

$$P_{mk} = \Big(g + \frac{g^{\text{sh}}}{2}\Big)|V_m|^2 - \frac{|V_k|}{\tau}|V_m|\big(g\cos(\theta_{km}-\sigma) - b\sin(\theta_{km}-\sigma)\big), \tag{2.10c}$$

$$Q_{mk} = -\Big(b + \frac{b^{\text{sh}}}{2}\Big)|V_m|^2 + \frac{|V_k|}{\tau}|V_m|\big(b\cos(\theta_{km}-\sigma) + g\sin(\theta_{km}-\sigma)\big). \tag{2.10d}$$

As we have seen, power can be specified as a function of the voltages at both ends of a branch. Therefore, when convenient, we will denote it as $S_{km}(|V_k|, |V_m|, \theta_k, \theta_m)$ and $S_{mk}(|V_m|, |V_k|, \theta_m, \theta_k)$ the complex power injected at each bus connected by branch km, respectively.

Power losses can also be computed taking into account the active power. The active power loss at branch km is defined as

$$\begin{aligned}\text{Loss}_{km} &= P_{km} + P_{mk} \\ &= \Big(g + \frac{g^{\text{sh}}}{2}\Big)\Big(\frac{|V_k|^2}{\tau^2} + |V_m|^2\Big) - 2\frac{|V_k|}{\tau}|V_m|g\cos(\theta_{km}-\sigma).\end{aligned} \tag{2.11}$$

Finally, we define the *apparent power* to be the magnitude (as complex number) of the power injected by k and m to branch km:

$$|S_{km}| = \sqrt{P_{km}^2 + Q_{km}^2} = ||(P_{km}, Q_{km})|| \tag{2.12a}$$

$$\text{and} \qquad |S_{mk}| = \sqrt{P_{mk}^2 + Q_{mk}^2} = ||(P_{mk}, Q_{mk})||, \tag{2.12b}$$

respectively. The apparent power has the same units as active or reactive power, however it is common to use the unit *voltampere* or *[VA]*. Moreover, the branch has a *capacity* $S_{km}^{\max}$ that limits the amount of apparent power.

2.1.2 Generation and Demand

Every generator in the system produces power that is injected into the network. We denote by

$$S_k^g = P_k^g + jQ_k^g \tag{2.13}$$

the amount of complex power produced by generator $k \in \mathcal{G}$, which must lie inside its capacities of generation. Depending on the physical capabilities of the generator, generated active power has as lower and upper bounds $P_k^{g,\min}$ and $P_k^{g,\max}$, respectively. Analogously, generated reactive power has as lower and upper bounds $Q_k^{g,\min}$ and $Q_k^{g,\max}$, respectively.

Electrical power is also consumed across buses of the network. It is common to say that the *load* of a bus is the amount of power that the bus demands, which is denoted by

$$S_k^d = P_k^d + jQ_k^d \tag{2.14}$$

for bus $k \in \mathcal{N}$.

According to [7], power systems are split into balancing authority areas that are in charge of maintaining the balance between load, generation, and interchange of power between the different balancing authority areas. This goal is accomplished through several processes, including

- Unit Commitment (UC): determines which generators will be on and off in an hourly based schedule, these generators are always ready to satisfy the forecasted demand;
- Economic Dispatch (ED): assigns the production level of the generators that were selected by the UC procedure, minimizing the cost, subject to meeting the demand and physical constraints. The ED algorithm is run every hour using the day-ahead load prediction, but also every 5-10 minutes, using the minute-ahead forecast;
- Primary Generation Control (PGC): performs on-line adjustments caused by generation outages, line tripping, demand fluctuations, or any spontaneous change in the system;

- Automatic Generation Control (AGC): returns the frequency —the angular speed of generator's rotors (see Section 2.3 for further details)— to its nominal value of 60 Hz after PGC actions; this is obtained by reassigning the generation output across different balancing authority areas; AGC updates commands every 2-4 seconds.

2.1.3 Operational Considerations

As described in [13, p. 327], operational consideration must be taken into account with the objective of understanding state and control variables of the system:

- For generators, the active power P_k^g and the voltage magnitude $|V_k|$ can be specified —by varying turbine power and generator field current.
- Complex load S_k^d at every bus is known (by previous estimation), thus, is considered as a parameter of the system.

We see then that active power can be set for all buses, however, this cannot be done independently, the sum of all the P_k^g's must equal the sum of active power loads and losses. One generator is designated to have unspecified active power generation, and this amount comes as the result of the system's active power balance. This specific bus is called a *slack bus* or *swing bus* (for which we will reserve the index $k = 0$), that instead of having specified P_0^g, both magnitude $|V_0|$ and angle θ_0 voltage values are indicated. Since power flow equations (2.10) always depend on the difference between the angle phases of two buses, by specifying the slack bus phase angle, we just set an angle reference for the rest of the buses.

In summary, we have three types of buses, depending on the specified variables:

1. A voltage V source. At the slack bus.
2. $P, |V|$ sources. At the other generators, also called *voltage control buses*.
3. P, Q sources. At the load buses.

2.2 Optimal Power Flow Problems

2.2.1 AC-OPF

The Alternative Current Optimal Power Flow (AC-OPF) problem is formulated in order to find an optimal solution of a power flow that minimizes the generated power subject to physical laws of energy conservation. AC-OPF was introduced in 1962 as the problem of economic dispatch [26].

Summarizing the previous sections, we have as parameters of the system:

- $V_k^{\min}$ and $V_k^{\max}$, voltage magnitude bounds for every bus $k \in \mathcal{N}$;
- $S_k^d = P_k^d + jQ_k^d$, the complex load at every bus $k \in \mathcal{N}$;
- $P_k^{g,\min}$ and $P_k^{g,\max}$, active power generation bounds for every generator $k \in \mathcal{G}$;
- $Q_k^{g,\min}$ and $Q_k^{g,\max}$, reactive power generation bounds for every generator $k \in \mathcal{G}$;
- $Y_{km} = \left[\begin{smallmatrix} y_{kk} & y_{km} \\ y_{mk} & y_{mm} \end{smallmatrix}\right]$, the branch admittance matrix, for every transmission line $km \in \mathcal{E}$;
- $S_{km}^{\max}$, the apparent power capacity for every branch $km \in \mathcal{E}$;
- $\theta_{km}^{\min}$ and $\theta_{km}^{\max}$, bounds of the voltage angle difference for every branch $km \in \mathcal{E}$;
- slack bus phase angle, usually specified as 0.

The control variables of the systems are:

- $|V_k|$, voltage magnitude at every generator $k \in \mathcal{G}$;
- θ_0, voltage phase angle of the slack bus;
- P_k^g, active power generation at every non-slack generator $k \in \mathcal{G}\backslash\{0\}$.

And the state variables are:

- $|V_k|$, voltage magnitude at every load bus $k \in \mathcal{N}\backslash\mathcal{G}$;

- θ_k, voltage phase angle at every non-slack bus $k \in \mathcal{N}\backslash\{0\}$;
- P_0^g, active power generation at the slack bus;
- Q_k^g, reactive power generation at every generator $k \in \mathcal{G}$;
- P_{km} and P_{mk}, active power injection at both ends of every branch $km \in \mathcal{G}$;
- Q_{km} and Q_{mk}, reactive power injection at both ends of every branch $km \in \mathcal{G}$.

The AC-OPF problem is then formulated as follows:

$$\begin{aligned}
\text{minimize} \quad & f(\{P_k^g\}_{k\in\mathcal{G}}) && \text{(2.15a)}\\
\text{subject to} \quad & \forall k \in \mathcal{G},\ (P_k^g - P_k^d) + j(Q_k^g - Q_k^d) = \sum_{km\in\delta(k)} (P_{km} + jQ_{km}) && \text{(2.15b)}\\
& P_k^{g,\min} \le P_k^g \le P_k^{\mathrm{g,max}}, \qquad Q_k^{g,\min} \le Q_k^g \le Q_k^{g,\max} && \text{(2.15c)}\\
& \forall k \in \mathcal{N}\backslash\mathcal{G},\ -(P_k^d + jQ_k^d) = \sum_{km\in\delta(k)} (P_{km} + jQ_{km}) && \text{(2.15d)}\\
& \forall km \in \mathcal{E},\ P_{km} + jQ_{km} = \mathrm{y}_{kk}|V_k|^2 + \mathrm{y}_{km}|V_k||V_m|e^{j(\theta_k-\theta_m)} && \text{(2.15e)}\\
& P_{mk} + jQ_{mk} = \mathrm{y}_{mm}|V_m|^2 + \mathrm{y}_{mk}|V_k||V_m|e^{-j(\theta_k-\theta_m)} && \text{(2.15f)}\\
& P_{km}^2 + Q_{km}^2 \le (S_{km}^{\max})^2, \quad P_{mk}^2 + Q_{mk}^2 \le (S_{km}^{\max})^2 && \text{(2.15g)}\\
& \theta_{km}^{\min} \le \theta_k - \theta_m \le \theta_{km}^{\max} && \text{(2.15h)}\\
& \forall k \in \mathcal{N},\ V_k^{\min} \le |V_k| \le V_k^{\max} && \text{(2.15i)}\\
& \theta_0 = 0. && \text{(2.15j)}
\end{aligned}$$

The objective function f depends on the active power generation at all generators, convex functions —typically quadratic— are commonly used so as to minimize the amount of power generated.

Constraints (2.15b) and (2.15d) describe the complex power balance at generators and load buses, respectively; generated power minus demand must equal the sum of power injected into the branches incident to the bus. Whereas constrain (2.15c) states the bounds for the power that a generator can produce.

Constrains (2.15e)-(2.15f) summarize the power equations that buses inject at their incident branches. These equations are explicitly stated in (2.10). Constraint (2.15g) limits the amount of apparent power that can be injected at each branch, while constraint (2.15h) states the bounds for the difference of voltage phase angles between two adjacent buses.

Finally, constraint (2.15i) established the bounds for the voltage magnitude and constraint (2.15j) sets the reference phase angle for the slack bus.

The AC-OPF problem formulated as in (2.15) is non-linear and non-convex; the source of non-linearity is mainly provided by the power flow equations (2.15e)-(2.15f), since they depends though trigonometric functions and products of the problem variables $\{|V_k|, \theta_k\}_{k\in\mathcal{N}}$. It is proved in [24] that finding feasible flows for the AC-OPF problem in general graphs is strongly NP-hard. However, numerical solvers can find approximate solutions for test networks with few thousand buses.

Approximations and relaxations of the AC-OPF problem are widely analyzed. The most simple and used one is the linear approximation, the so-called DC approximation, that will be described in the next section. [55] and [54] provide several relaxations using Second Order Cone Programming (SOCP) and Semidefinite Programming (SDP).

2.2.2 DC-OPF

In this section we describe a linear approximation of the AC-OPF problem (2.15). Given its simplicity and fast implementation, the DC-OPF is the problem that the economic dispatch procedure solves every 5-10 minutes.

Under normal conditions, the following observations are made in order to simplify constraints (2.15e)-(2.15f):

- for every bus $k \in \mathcal{N}$, $|V_k| \approx 1$ (using per-unit system),
- for every branch $km \in \mathcal{E}$,
  - $\theta_{km} = \theta_k - \theta_m$ is small so that $\sin(\theta_{km}) \approx \theta_{km}$ and $\cos(\theta_{km}) \approx 1$,

- resistance is much smaller that reactance, meaning that $r_{km} \ll x_{km}$. In consequence, $g_{km} \approx 0$ and $b_{km} \approx -1/x_{km}$,
- shunt admittance and transformer effects are ignored, that is, $y_{km}^{\text{sh}} = 0$ and $N_{km} = 1$.

With these observations, expressions (2.10) for active and reactive power become:

$$P_{km} = -b_{km}(\theta_k - \theta_m) = \frac{\theta_k - \theta_m}{x_{km}}, \qquad Q_{km} = 0, \tag{2.16a}$$

$$P_{mk} = b_{km}(\theta_k - \theta_m) = -\frac{\theta_k - \theta_m}{x_{km}}, \qquad Q_{mk} = 0. \tag{2.16b}$$

In other words, in the DC approximation the transmission lines are lossless ($P_{km} + P_{mk} = 0$) and carry only active power and, therefore, the reactive power —and imaginary part of variables and equations— is completely ignored. The balance equations (2.15b) and (2.15d) are simplified to

$$\forall k \in \mathcal{N}, \quad P_k^g - P_k^d = \sum_{km \in \delta(k)} \frac{\theta_k - \theta_m}{x_{km}}, \tag{2.17}$$

where we define as $P_k^g = 0$ when $k \notin \mathcal{G}$. For each choice of nonnegative values P_k^g and P_k^d such that $\sum_{k\in\mathcal{N}} P_k^g = \sum_{k\in\mathcal{N}} P_k^g$, the system (2.16)-(2.17) has a unique solution in the P_{km} (Lemma 1.1 in [23]).

An equivalent formulation is given by defining the vectors $\theta = (\theta_k : k \in \mathcal{N})$, $\mathcal{P}^g = (P_k^g : k \in \mathcal{N})$ and $\mathcal{P}^d = (P_k^d : k \in \mathcal{N})$ and the $|\mathcal{N}| \times |\mathcal{N}|$ *bus susceptance* matrix $\mathcal{B}$ as

$$\forall k \in \mathcal{N},\ \mathcal{B}_{kk} = \sum_{km \in \delta(k)} \frac{1}{x_{km}}, \qquad \forall km \in \mathcal{E},\ \mathcal{B}_{km} = \mathcal{B}_{mk} = -\frac{1}{x_{km}}, \tag{2.18}$$

and $\mathcal{B}_{km} = 0$, otherwise, we can restate the power balance system (2.17) as the $|\mathcal{N}| \times |\mathcal{N}|$ linear system

$$\mathcal{P}^g - \mathcal{P}^d = \mathcal{B}\,\theta. \tag{2.19}$$

Because of the construction of $\mathcal{B}$, the set of equations of the linear system (2.19) is linearly dependent and is feasible only if $\sum_{k\in\mathcal{N}}(P_k^g - P_k^d) = 0$. By adding the equation $\theta_0 = 0$ that sets the phase angle for the slack bus, the system becomes linearly independent

when the underlying network is connected. In this case, the solution to system (2.19) has a unique solution, of the form

$$\theta = \breve{\mathcal{B}}_0(\mathcal{P}^g - \mathcal{P}^d) \tag{2.20}$$

where $\breve{\mathcal{B}}_0$ is an appropriate pseudo-inverse of $\mathcal{B}$ which depends on the choice of the reference (slack) bus.

As in the case of the AC formulation, P_k^g for $k \in \mathcal{G}\backslash\{0\}$ are control variables, and the generation at the slack bus is set to be $P_0^g = \sum_{k\in\mathcal{N}} P_k^d - \sum_{k\in\mathcal{G}\backslash\{0\}} P_k^g$.

The DC-OPF problem is formulated as the following linearly constrained system:

$$\begin{aligned}
\text{minimize} \quad & f(\mathcal{P}^g) && \text{(2.21a)}\\
\text{subject to} \quad & \mathcal{P}^g - \mathcal{P}^d = \mathcal{B}\,\theta, \quad \theta_0 = 0 && \text{(2.21b)}\\
& \forall k \in \mathcal{G},\ P_k^{g,\min} \le P_k^g \le P_k^{g,\max} && \text{(2.21c)}\\
& \forall km \in \mathcal{E},\ -S_{km}^{\max} \le (\theta_k - \theta_m)/x_{km} \le S_{km}^{\max} && \text{(2.21d)}\\
& \theta_{km}^{\min} \le \theta_k - \theta_m \le \theta_{km}^{\max}. && \text{(2.21e)}
\end{aligned}$$

If f is a convex function, then program (2.21) is a convex optimization problem.

2.2.3 Numerical Solutions

There exist solvers that find the solution (or an approximate optimum when the problem is difficult) to the optimization problems stated above; either in their OPF form —that finds the optimum flow— or in their PF version, that just finds an (approximate) feasible solution to a power flow.

MATPOWER [101] is a MATLAB toolbox widely used in the literature that solves optimization power flow instances and has a library of cases —usually known as IEEE test cases— with different number of buses, from 10 to 10,000, and serves to test algorithms, theory and their implementation.

2.3 Time-Varying Elements

In this section we describe how the elements detailed in the previous sections evolve as time varies. The optimization problems stated in Section 2.2 that find optimal power flows are solved considering one particular instant of time, however, as generation must match demand —plus losses— at every moment, variables and parameters evolve dynamically with time.

We introduce a time dependent argument to the following variables and parameters:

- voltage $V_k(t) = |V_k(t)|e^{j\theta_k(t)}$,
- demand $S_k^d(t) = P_k^d(t) + jQ_k^d(t)$ and
- generation $S_k^g(t) = P_k^g(t) + jQ_k^g(t)$.

In consequence, since current and power depend on the voltages, they are also time dependent, denoted as $I_{km}(t)$ and $S_{km}(t) = P_{km}(t) + jQ_{km}(t)$, respectively.

In steady state, voltages and currents are sinusoidal functions of time. That is, voltage of bus k can be written as

$$V_k(t) \;=\; |V_k(t)|e^{j\theta_k(t)} \qquad \text{with} \quad \theta_k(t) = \omega_k(t) \cdot t + \delta_k(t), \tag{2.22}$$

where $\omega_k(t)$ is the voltage angular frequency —or, simply, frequency— and $\delta_k(t)$ the voltage phase angle referenced to $\omega_k(t) \cdot t$. For steady state analyses it is usual to consider that $\omega_k(t)$ and $\delta_k(t)$ slowly change through time or do not change at all and, in consequence, they are assumed to be constant functions. Moreover, the frequency is considered equal for all buses at $\omega_k(t) = \omega_0 = 60$ Hz. This is explained by the fact the power injected into the system is produced by generators that spin at approximately 60 Hz, see Figure 2.2 (detailed description of generators can be found in [13]).

The grid has been built in such a way that the generators' rotor must synchronously spin at a regulated frequency —60 Hz— to avoid mechanical resonances. However, in practice, the frequency deviates from 60 Hz by small amounts, that is, $\omega_k(t) = 60 + \epsilon_k(t)$

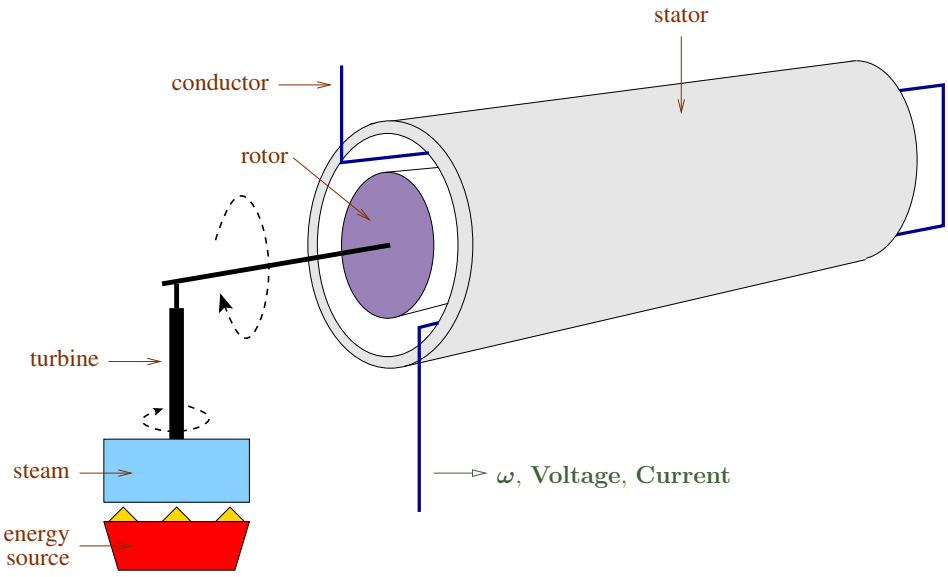

Figure 2.2: Turbine-Generator Model. A turbine converts different type of energy —such as steam or burnt natural gas— into rotation. The turbine shaft is connected to a generator shaft. A generator typically consists in a rotor (rotating part) that spins at a frequency $\omega(t)$ inside a stator (static part). The rotation of the rotor creates a magnetic field that induces current into the conductor at frequency $\omega(t)$.

with $|\epsilon_k(t)|$ small. In normal operation, the frequency is controlled such that $|\epsilon_k(t)| < 0.05$ Hz. Nevertheless, frequency can deviate as much as 5% for short periods without causing load shedding, generator tripping, damaged equipment, or threatening the stability of the system [53].

Fluctuations of frequency are caused by generation outages, tripping lines, intermittent generation or fluctuations of demand. When demand suddenly increases, the inertia maintained by spinning rotors —mechanical energy— is transferred into electrical energy to overcome the power imbalance, therefore, rotors slow down and frequency decreases. Figure 2.3 shows a contingency where 2,600 *[MW]* of generation is lost, as a consequence, the frequency drops almost instantaneously but it is afterwards recovered by automated measures (Automatic Generation Control).

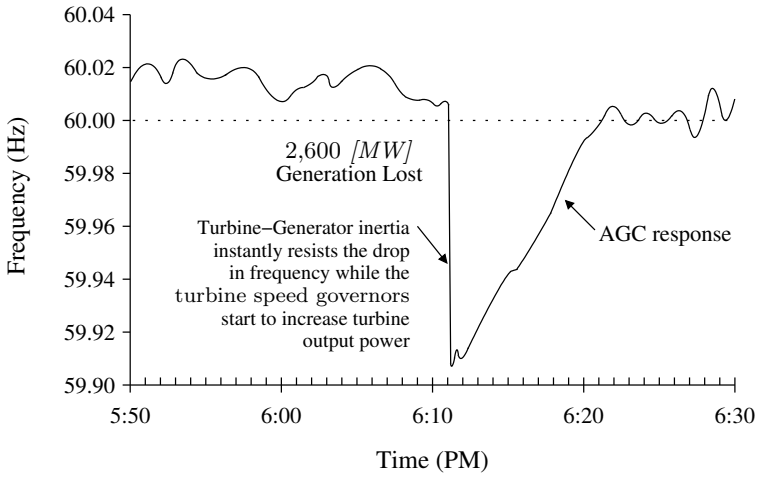

Figure 2.3: Contingency reserves rebalance the system after the sudden loss of generation [53].

2.4 State Estimation

State estimation (SE) has served as a tool to supervise and control the functioning of a power system. Thanks to SE, the grid's controller can take on-line decisions with the objective of optimizing processes, adjusting generation, and detecting bad data [46].

As we have seen, the state of the system can be described by their state variables (voltages at all buses) together with the physical characteristics of the grid. A set of sensors located through the network make measurements of some of these variables, and send this information to the control center. A computational tool uses this data to perform SE in the following way: first, it undertakes an observability analysis that concludes if the received data is enough to get a good state estimation of the system; next, bad measurements are identified and eliminated of the data set; finally, an optimization problem is solved to extrapolate the measured and trusted sampling to the rest of variables that cannot be observed.

The most advanced sensors are the PMUs (phasor measurement units) that are located on branches next to a bus. They are synchronized by a GPS clock and report, 30 times a second, data that includes time of the measurement, frequency $\omega_k(t)$, voltage phase angle $\theta_k(t)$ and voltage magnitude $|V_k(t)|$ of the adjacent bus, status of the branch

(active or inactive), complex current $I_{km}(t)$ of the branch, and other data.

Traditional methods for determining the state variables are solved by using the weighted least-squares (WLS) method. For that purpose, if z is a vector of observable quantities, then the solution x of the system

$$z \;=\; h(x) + e, \tag{2.23}$$

gives us the state variables (voltage phase angles and voltage magnitudes) of the network, where h is a set of nonlinear functions of the state variables and e is a zero-mean Gaussian measurement noise vector with covariance matrix C.

Then, the SE problem finds

$$\hat{x} \;=\; \arg\min_{x} [z - h(x)]^{\top} W^{-1} [z - h(x)],$$

where the weighting matrix W is usually taken diagonal and it is related with the noise covariance matrix C.

Chapter 3

Cyber-Physical Attacks

In this section we introduce the notion of a cyber-physical attack on power systems. We also model a complex attack that we demonstrate to be feasible on large networks. Finally, we describe defense mechanisms against these and more general attacks.

3.1 Introduction

Cyber-physical attacks have been widely analyzed in the last decade. The main purpose of this study is to investigate complex weaknesses of the grid and elaborate mechanisms to strengthen the security of the system.

The *cyber* adjective of the attacks correspond to the notion that sensor measurements used to estimate the state of the system are perturbed or deleted. Voltage and current throughout the network are periodically measured with sensors (PMUs, RTUs) and sent to the *control center* of the grid. Using this data, the control center performs state estimation, that is, using incomplete knowledge of the all the variables of the system, unknown parameters and variables are estimated (typically by minimizing square errors) with the objective of supervising normal operation and plan-ahead future control actions (see Section 2.4 for details). If the information received by the control center is modified, actions that can be damaging to the network stability might be taken.

A *physical* attack consists on a physical alteration of the system or a portion of it, such as changes in generation or demand (e.g. by shutting down a generator or coordinated activation of many air conditioners), line tripping, alteration of physical properties of lines or renewable energy generation, among others.

Thus, a cyber-physical attack has both events happening together and it is typically performed by an adversarial agent. One of the iconic examples of a cyber-physical attack that has recently taken place is the attack on the Ukrainian electrical system on 23 December 2015 where attackers gained access to the data acquisition system, were able to delete and modify data, and also changed the grid topology by activating circuit breakers [34]. In consequence, 225,000 customers lost power for a period of 1 to 6 hours.

In this work, we propose defensive techniques to be deployed when a high-fidelity attack on a power grid is suspected. The attack is assumed to be *partial* in the sense that only a subset of buses and lines are attacked, but this subset is unknown by the system controller. These techniques involve two ideas:

(a) using network resources to randomly change power flow quantities, especially voltages and, in particular

(b) changing the covariance structure of e.g. voltages in a manner unpredictable by the attacker. The specific version of this idea that we analyze introduces a low-rank adjustment to the covariance of phase angles.

A precise definition of the attack is given in Section 3.3. These defensive techniques focus on the phase following the initial attack, and aim to expose inconsistencies in the modified sensor data stream which is output by the attacker. We describe conditions under which the defense succeeds in discovering the *boundary* of the attacked zone.

We justify such defenses by pointing out that the possibility of dangerous "cyber-physical" attacks of high-fidelity and with sparse signatures has already been indicated in the literature (see Section 3.2). The data component of the attacks is designed to pass a stringent test, namely that the falsified data satisfies the full AC power flow equations

at every bus and line. The data attack is coordinated with a physical attack encompassing various types (in particular, line tripping, or load modifications as considered in this paper) that results in a dangerous system condition, e.g. a line overload. The data modification hides this overload, with the result that sensor data received by operators is both unimpeachable and portrays safe system operation. We term these attacks "ideal" because, while sparse, they do assume technical sophistication and the ability to coordinate physical action and computation. Sparsity is a goal for the attacker because it increases the likelihood of undetectability long enough for the overload to lead to line tripping (typically several minutes). Putting aside the actual feasibility of such attacks, the computational challenge is significant.

We will start by giving a summary of related work. Then, we develop a new optimization procedure that successfully computes cyber-physical attacks on large transmission systems with thousands of buses. Finally, we will focus on the defensive mechanisms against ideal cyber-physical attacks.

3.2 Previous Work

Early work on full observability of grid sensor data [66] and bad data detection has been done in [11,30,61,68,69]. These works have developed techniques to detect bad measurements and remove them from the data set. The principal tool that they use is based on the fact that the objective function minimized (e.g. squared errors) during state estimation becomes significantly high when *arbitrary* bad measurements are present. As shown in [64], these techniques do not work well when an adversary —having enough knowledge of the system— introduces false data, since it could make the injected information look coherent with some state of the grid. [64] was the first to introduce the *false data injection* (FDI) attacks, where they use the DC power flow model for attack analysis, arguing that it is less accurate, but more simpler and robust than the AC model [90].

False Data Injection Attacks

As explained before, this kind of attacks rely on the modification of the measurements received by the control center affecting a subset of observable variables. [64, 87] study cyber attacks where an attack vector a is added by the adversary to the observed vector z in the system of equations (2.23) to create a malicious measurement vector $z^a = z + a$ (the vector a has non-zero entries only in the components related with the sensors that the attacker has access to). Ideally, the attacker expects that a vector x^a minimizes the errors of the system, with low weighted least-squares error, so that x^a appears to be a reasonable state variables vector. They describe results that apply when the attacker has limited access to certain sensors or when he has a budget of how large the changes can be. [51] also proves the existence of attack vectors that belong to a subspace of the observable vector domain, which becomes relevant when the attacker does not have full observability of the sensor's readings. In [47], this class of attacks are extended to AC state estimation, creating unobservable attacks on specific sub-graphs of the network.

In [35] and [73] consider an attacker with no information about the grid topology. Under certain assumptions, the attacker can learn the topology from the power flow observations an still launch undetectable cyber attacks.

In [50] the authors study an FDI attack including topology modification, in particular, it is consistent with a solution where branches are either not working but in reality they are, and viceversa. The paper considers cases when either the attacker has full information (observability) of the network or just some local knowledge. Meanwhile, the attacks proposed in [59] based in the AC SE can lead to generation re-dispatch and line overflows. The work presented in [63, 95] analyzes FDI attacks that cause a load redistribution in the SE using the DC power flow model.

Authors in [32] and [52] propose different algorithms to establish a minimum-cost set of sensors should be secured in order to prevent the kind of cyber attacks described above. Whereas [91] defines two metrics to quantify the importance of substations (how many different attack vectors can involve a particular sensor) and the cost of the attacks

on specific measurements (how many sensors are needed to be perturbed in order to alter the measurement); and proposes strategies improve the system security with respect to these metrics.

Cyber-Physical Attacks

The study in [80] considers a cyber-physical attack under the DC model where an adversary attacks a zone H of nodes by disconnecting some lines (without disconnecting the graph) and obstructing the measurements within H. (They assume that after the attack there is no edge $\{i, j\} \in E_H$ such that $\theta'_i = \theta'_j$.) As a first step, they show that if there is a matching between nodes outside H and inside H that covers the nodes in H, the phase angles can be recovered almost surely (in the sense that for very specific combination of reactance values of the edges in the matching the phase angles could not be recovered). Secondly, the authors prove that the support of a linear system solution coincides with the set of lines that failed (disconnected lines), and this solution is unique if and only if the subgraph induced by nodes in the attacked zone H is acyclic. Also, if H is a cycle and strictly less than half of its edges are disconnected, the support of a linear program gives the set of failed lines. Other sufficient conditions are developed for planar graphs. They extend these results in the presence of noisy data.

In [82] the authors propose an algorithm to estimate the phase angle after attacks (that trip "any number" of lines and obstruct information from the attacked zone) using a convex relaxation of the estimation for DC power flows (relaxing equalities to second order cone inequalities), no specific/theoretical use of AC power flow equations.

The work presented in [83] also study cyber-physical attacks affecting the physical infrastructure and the SCADA (Supervisory Control And Data Acquisition) system of the network. This work is an extension of the results showed in [80] but using the AC power flow model: an attacker disconnects lines from an area and obstructs the information measured in that area, here generation and load remain the same after the attack for all nodes. The paper introduces the EXPOSE Algorithm to detect line failures using the AC

power flow equations, the algorithm is independent of the size of the grid and the number of line failures. Together with [82], the authors claim to have the only methods that scale with the size of the grid to detect any number of line failures with the AC power flow equations. Using a similar idea than in [80], they show that the complex voltages of the attacked zone can be recovered almost surely if there is a matching that covers the attacked zone nodes by solving a linear system. Also, the authors state the same result than in [80] when the attacked zone is a cycle. In more general cases, they propose the EXPOSE Algorithm that solves a convex minimization problem that assuming that the voltage magnitude of the attacked buses remain constant after the attack.

The authors in [96] study cyber-physical attacks with the following characteristics: an intelligent attacker has access and gets knowledge of the network (topology, operation costs, historical data, observable measurements), the attacker performs an attack that trips a specific transmission line with the objective of overload another chosen line, this physical attack is coordinated with a cyber attack that masks the tripped branch and avoids its detection. The proposed strategy has two steps: first, the detection of a set of buses able to attack and a target line to trip; and second, the computation of an attack vector to inject as false data. As in [60], the authors compute first a DC attack vector by solving a linear optimization problem, and then, using AC state estimation corresponding to the DC attack vector, they compute a more accurate AC attack vector. They test their results in the IEEE 24-bus test case. An attack of similar characteristics using the DC power flow model is proposed in [33] together with countermeasures to detect the intrusion: the availability of trusted PMUs and tracking of power system equivalent impedance. Tests in the IEEE 9-bus, 14-bus, 30-bus, 118-bus, and 300-bus test power systems are shown.

[58] also study cyber-physical attacks that, in this case, the physical attack preserves the grid topology but redistributes loads (see also [63, 95]) within certain budget across the network and hides it through manipulated injection of data, under the DC model. The paper bases the analysis on the computation of the *generation shift factors* that

capture sensitivity of power flow changes on lines when bus injections are altered. Two-stage mixed-integer linearized optimization problems are proposed to obtain the set of lines and buses that will be affected by the attack. Given the dimensionality of the problem, the authors show results of coordinated cyber-physical attacks on small IEEE test systems with 14 and 118 buses.

Other Related Literature

The work in [28] proposes a game theory model in where an attacker and a defender of the network have limited budget that allow them to attack targeted elements of the grid (trip a line, break a transformer) o to secure them, respectively. They implement an algorithm and show that when defenders deploy an strategy before an attack is initiated, the loss can be predictable and limited to a minimum level.

Other works have focused on the system vulnerability to line outage, see e.g. [14, 18, 29, 39, 72, 85, 86, 97, 98, 100]. The objective of these studies is to identify from the data when line tripping has happened, but also analyze the consequence that they could caused, for example, cascade line failures. In particular, [72] uses a bilevel mixed integer nonlinear programming to formulate the problem of finding a small set of lines whose removal from the network would cause severe blackouts.

3.3 Attack Model

We will now describe and formulate the cyber-physical attacks that we considerate, combining physical disruption and data intrusion. We will show that the numerical solution to this problem scales well to systems with thousands of buses, with running times in the tens of seconds or less on a standard computer. These experiments do not prove that the attacks, though sparse, could easily be executed. Rather they show that the attack computation is tractable, thus providing added justification for studying sophisticated and scalable defense mechanisms.

Attack models considered in prior work allow the attacker different capabilities. Regardless of the model, Template 1 given below (similar to one in [96]) broadly outlines the structure of an AC-undetectable attack. We use the term "initial" to indicate that the attack comprises actions taken at one point in time. Later we will discuss a "follow-up" phase that follows the initial attack.

Template 1. *Initial Attack*

(a) It is assumed that at each bus k there is a sensor measuring voltage at k and current at each line $km \in \delta(k)$.

(b) The attacker has selected a (sparse) subset $\mathcal{A}$ of buses, as well as a target line uv within $\mathcal{A}$ that will be overloaded.

(c) For any bus $k \in \mathcal{A}$, the attacker can modify data provided by a sensor located at k.

(d) The attacker's physical actions are of two types. First, the attacker can modify *loads* at buses in $\mathcal{A}$. Additionally the attacker can disconnect lines with both ends in $\mathcal{A}$.

(e) Actions (c)-(d) are performed in a single step.

(f) The data received by the control center satisfies complete fidelity as per AC power flow equations (2.7)-(2.10) and shows all system limits being satisfied, while in actuality line uv is overloaded.

(g) When the attack includes load changes, secondary response (i.e. AGC response) is taken into account by the attacker.

Conditions (a), (f) and (g) amount to a strong form of undetectability. Nevertheless, we provide examples of large scale systems that are susceptible to attacks of the form

(a)-(g). Note that we allow loads to be modified, but not generation. In our numerical examples we enforce that $\mathcal{G} \cap \mathcal{A} = \emptyset$, out of a perception that generator sites are more carefully protected.

Point (e) requires some discussion. Immediately following any physical modification to a system, we can expect a change in voltages (magnitudes and phase angles) and even to system frequency, the latter especially when net loads are changed. More properly, system dynamics will undergo a change. Understanding the precise nature of that change is a substantial computational task. The current state-of-the-art involves a numerical simulation that alternates between simulation of true dynamical behavior at generators (the so-called swing equation) with AC power flow updates. This combined computation will typically run much slower than the actual dynamics, and assumes correct knowledge of the underlying transmission system. Under adversarial attack that e.g. modifies the topology, the rapid success of such a computational approach to identifying the current grid state seems uncertain. And once action (d) (i.e. modification of sensor data) is taken a completely falsified, and consistent view of the system is being presented.

We next present conditions that we will impose so as to guarantee undetectability. *True* data will be the true physical data. In contrast, *reported* data is that which is actually received by the control center and includes the attacker's modifications. The true data will be given by the (voltage, current) pair of vectors $(V^{\mathrm{T}}, I^{\mathrm{T}})$ whereas the reported data will be given by $(V^{\mathrm{R}}, I^{\mathrm{R}})$.

An important requirement for the reported data is

$$\textbf{current-voltage consistency:} \qquad \begin{pmatrix} I^{\mathrm{R}}_{km} \\ I^{\mathrm{R}}_{mk} \end{pmatrix} = Y_{km} \begin{pmatrix} V^{\mathrm{R}}_{k} \\ V^{\mathrm{R}}_{m} \end{pmatrix}, \tag{3.1}$$

i.e. equation (2.8), should be satisfied at all branches. This condition will be enforced in the computation given below in an indirect fashion (also see [67] for a different use of this requirement). In general, of course, an attacker might only seek approximate consistency, using ambient noise to hide errors. Additionally to (3.1):

(s.1) On a bus $k \notin \mathcal{A}$ the true and reported data agree (no data modification outside $\mathcal{A}$,

by definition).

(s.2) At a boundary bus[1] $k \in \partial\mathcal{A}$ the attacker is constrained by the condition $V_k^{\mathrm{R}} = V_k^{\mathrm{T}}$. This condition is applied to avoid attack detection, given (a) and the second equation in (3.1) applied to a line km where $m \notin \mathcal{A}$.

(s.3) On buses $k \in \mathcal{A}\backslash\partial\mathcal{A}$ we may have $V_k^{\mathrm{R}} \neq V_k^{\mathrm{T}}$ and on lines with at least one end in $\mathcal{A}\backslash\partial\mathcal{A}$ the true and reported currents may also differ.

(s.4) The reported voltages and currents must be consistent with meaningful (complex) power injections. Specifically, consider a bus k. Then $\sum_{km\in\delta(k)} V_k^{\mathrm{R}} I_{km}^{\mathrm{R}*}$ equals the power injected into the system at bus k, according to the reported data. If $k \notin \mathcal{A}$ by definition (of reported and true data) this sum equals $\sum_{km\in\delta(k)} V_k^{\mathrm{T}} I_{km}^{\mathrm{T}*}$ which is the true power injected by bus k. On the other hand if $k \in \mathcal{A}$ the sum may differ from the true injection at k.

(s.5) If the attack causes a net change in the sum of loads, the resulting AGC-mandated change in generator output must be taken into account.

We call condition (s.4) **power-injection consistency**, that is:

$$\sum_{km\in\delta(k)} V_k^{\mathrm{R}} I_{km}^{\mathrm{R}*} = \text{net injection at } k \quad \forall k, \tag{3.2}$$

where "net injection" is the reported net injection on buses in $\mathcal{A}$ and the true net injection for buses not in $\mathcal{A}$.

Subject to these requirements, the attacker seeks to create a (true) line overload on uv with both ends in $\mathcal{A}$, while the reported data shows safe system operation (voltage, angle, and power flow limits are satisfied). In other words, for $\eta > 1$ (e.g. $\eta = 1.2$ or

[1]Recall that $k \in \mathcal{A}$ is a boundary bus if it has a neighbor outside of $\mathcal{A}$.

$\eta = 1.5$), the attacker would try to satisfy

$$(P_{uv}^{\mathrm{T}})^2 + (Q_{uv}^{\mathrm{T}})^2 = (S_{uv}^{\mathrm{T}})^2 > (\eta S_{uv}^{\max})^2 \tag{3.3}$$

$$\forall km \in \mathcal{E}, \quad (P_{km}^{\mathrm{R}})^2 + (Q_{km}^{\mathrm{R}})^2 = (S_{km}^{\mathrm{R}})^2 \leq (S_{km}^{\max})^2 \tag{3.4}$$

$$\theta_{km}^{\min} \leq \theta_k^{\mathrm{R}} - \theta_m^{\mathrm{R}} \leq \theta_{km}^{\max} \tag{3.5}$$

$$\forall k \in \mathcal{N}, \qquad V_k^{\min} \leq |V_k^{\mathrm{R}}| \leq V_k^{\max} \tag{3.6}$$

where $(P_{km}^{\mathrm{T}}, Q_{km}^{\mathrm{T}})$ and $(P_{km}^{\mathrm{R}}, Q_{km}^{\mathrm{R}})$ are the true and reported flows given by the expressions in (2.10) using $\{V_k^{\mathrm{T}}\}_k$ and $\{V_k^{\mathrm{R}}\}_k$, respectively. In the next section we present a mathematical formulation for this problem.

3.3.1 Formulation

As input to the problem we have a set $\mathcal{A} \subset \mathcal{N} \backslash \mathcal{G}$ of buses, a set of lines $\mathcal{L}$ to be disconnected, all with both ends in $\mathcal{A}$, and a line $uv \notin \mathcal{L}$ with both ends in $\mathcal{A}$. Write $\mathcal{A}^C = \mathcal{N} \setminus \mathcal{A}$. Let $(\hat{S}_k^g = \hat{P}_k^g + j\hat{Q}_k^g)_{k \in \mathcal{N}}$ and $(\hat{S}_k^d = \hat{P}_k^d + j\hat{Q}_k^d)_{k \in \mathcal{N}}$ be the complex power generation and loads, respectively, at the time of the attack. We assume that the attacker observes all these quantities.

The initial attack problem uses the following real-valued variables, where "T" indicates true and "R", reported:

- $\forall$ bus $k \in \mathcal{N}$: $|V_k^{\mathrm{T}}|, \theta_k^{\mathrm{T}}$ denote the true voltage magnitude and voltage angle;
- $\forall$ bus $k \in \mathcal{N}$: $|V_k^{\mathrm{R}}|, \theta_k^{\mathrm{R}}$ denote the reported voltage magnitude and voltage angle;
- $\forall$ bus $k \in \mathcal{A}$: $P_k^{d,\mathrm{T}}, Q_k^{d,\mathrm{T}}$ denote the active and reactive true loads in $\mathcal{A}$;
- $\forall$ bus $k \in \mathcal{A}$: $P_k^{d,\mathrm{R}}, Q_k^{d,\mathrm{R}}$ denote the active and reactive reported loads in $\mathcal{A}$;
- $\forall$ bus $k \in \mathcal{R}$: P_k^g, Q_k^g denote active and reactive generation at participating buses;
- $\forall$ line $km \in \mathcal{E} \backslash \mathcal{L}$: $P_{km}^{\mathrm{T}}, Q_{km}^{\mathrm{T}}$ denotes the active and reactive true power flows;
- $\forall$ line $km \in \mathcal{E}$: $P_{km}^{\mathrm{R}}, Q_{km}^{\mathrm{R}}$ denotes the active and reactive reported power flows;

- Δ denotes the net change in active power generation.

Therefore, the problem is formulated as follows:

$$\text{maximize} \quad (P_{uv}^{\mathrm{T}})^2 + (Q_{uv}^{\mathrm{T}})^2 \tag{3.7a}$$

$$\text{subject to} \quad \forall k \in \mathcal{A}^C \cup \partial\mathcal{A}, \quad |V_k^{\mathrm{T}}| = |V_k^{\mathrm{R}}|, \quad \theta_k^{\mathrm{T}} = \theta_k^{\mathrm{R}} \tag{3.7b}$$

$$\forall k \in \mathcal{A}, \quad -(P_k^{d,\mathrm{R}} + jQ_k^{d,\mathrm{R}}) = \sum_{km \in \delta(k)} (P_{km}^{\mathrm{R}} + jQ_{km}^{\mathrm{R}}), \tag{3.7c}$$

$$-(P_k^{d,\mathrm{T}} + jQ_k^{d,\mathrm{T}}) = \sum_{km \in \delta(k)\setminus\mathcal{L}} (P_{km}^{\mathrm{T}} + jQ_{km}^{\mathrm{T}}), \tag{3.7d}$$

$$P_k^{d,\mathrm{R}} \geq 0, \; P_k^{d,\mathrm{T}} \geq 0 \tag{3.7e}$$

$$\forall k \in \mathcal{A}^C \setminus \mathcal{R}, \quad \hat{P}_k^g - \hat{P}_k^d + j(\hat{Q}_k^g - \hat{Q}_k^d) = \sum_{km \in \delta(k)} (P_{km}^{\mathrm{T}} + jQ_{km}^{\mathrm{T}}) \tag{3.7f}$$

$$\forall k \in \mathcal{R}, \quad P_k^g - \hat{P}_k^d + j(Q_k^g - \hat{Q}_k^d) = \sum_{km \in \delta(k)} (P_{km}^{\mathrm{T}} + jQ_{km}^{\mathrm{T}}) \tag{3.7g}$$

$$P_k^g - \hat{P}_k^g = \alpha_k \Delta \tag{3.7h}$$

$$\forall k \in \mathcal{G}, \quad P_k^{g,min} \leq P_k^g \leq P_k^{g,\max}, \; Q_k^{g,min} \leq Q_k^g \leq Q_k^{g,\max} \tag{3.7i}$$

$$\forall k \in \mathcal{N}, \quad V_k^{min} \leq |V_k^{\mathrm{T}}|, |V_k^{\mathrm{R}}| \leq V_k^{\max} \tag{3.7j}$$

$$\forall km \in \mathcal{E}, \quad |\theta_k^{\mathrm{R}} - \theta_m^{\mathrm{R}}| \leq \theta_{km}^{\max}; \; |\theta_k^{\mathrm{T}} - \theta_m^{\mathrm{T}}| \leq \theta_{km}^{\max} \text{ if } km \notin \mathcal{L}, \tag{3.7k}$$

$$\max\{ \|(P_{km}^{\mathrm{R}}, Q_{km}^{\mathrm{R}})\|, \|(P_{mk}^{\mathrm{R}}, Q_{mk}^{\mathrm{R}})\| \} \leq S_{km}^{\max}, \tag{3.7l}$$

$$P_{km}^{\mathrm{T}} + jQ_{km}^{\mathrm{T}} = S_{km}(|V_k^{\mathrm{T}}|, |V_m^{\mathrm{T}}|, \theta_k^{\mathrm{T}}, \theta_m^{\mathrm{T}}) \quad km \notin \mathcal{L} \tag{3.7m}$$

$$P_{mk}^{\mathrm{T}} + jQ_{mk}^{\mathrm{T}} = S_{mk}(|V_m^{\mathrm{T}}|, |V_k^{\mathrm{T}}|, \theta_m^{\mathrm{T}}, \theta_k^{\mathrm{T}}) \quad km \notin \mathcal{L} \tag{3.7n}$$

$$P_{km}^{\mathrm{R}} + jQ_{km}^{\mathrm{R}} = S_{km}(|V_k^{\mathrm{R}}|, |V_m^{\mathrm{R}}|, \theta_k^{\mathrm{R}}, \theta_m^{\mathrm{R}}) \tag{3.7o}$$

$$P_{mk}^{\mathrm{R}} + jQ_{mk}^{\mathrm{R}} = S_{mk}(|V_m^{\mathrm{R}}|, |V_k^{\mathrm{R}}|, \theta_m^{\mathrm{R}}, \theta_k^{\mathrm{R}}) \tag{3.7p}$$

In this formulation, power flows are represented through the quadratic functions S_{km}, S_{mk} (see equations (2.9)-(2.10)) which appear in the formulation as (3.7m)-(3.7p). Note that we include voltage variables but no current variables. However, having solved the above optimization problem, the attacker reports, for each line km with both ends

in $\mathcal{A}$, a current pair $(I^{\mathrm{R}}_{km}, I^{\mathrm{R}}_{mk})$ computed using the formula

$$\begin{pmatrix} I^{\mathrm{R}}_{km} \\ I^{\mathrm{R}}_{mk} \end{pmatrix} = Y_{km} \begin{pmatrix} |V^{\mathrm{R}}_k| e^{j\theta^{\mathrm{R}}_k} \\ |V^{\mathrm{R}}_m| e^{j\theta^{\mathrm{R}}_m} \end{pmatrix},$$

thereby attaining current-voltage consistency (3.1). Note that if either $k \in \partial\mathcal{A}$ or $m \in \partial\mathcal{A}$ the true and reported voltage values are identical —see Lemma 3 below.

Lemma 1. *Consider a feasible solution to problem (3.7). Let* H *denote either* T *or* R *(i.e. true or reported). Then the voltages $|V^{\mathrm{H}}_k| e^{j\theta^{\mathrm{H}}_k}$ for all $k \in \mathcal{N}$ yield a solution to the power flow problem where*

(1) Bus k has load $P^{d,\mathrm{H}}_k + jQ^{d,\mathrm{H}}_k$ for $k \in \mathcal{A}$ and $\hat{P}^d_k + j\hat{Q}^d_k$ if $k \in \mathcal{A}^C$.

(2) Bus $k \in \mathcal{G}$ has generation $P^g_k + jQ^g_k$ if $k \in \mathcal{R}$ and $\hat{P}^g_k + j\hat{Q}^g_k$ if $k \in \mathcal{G} \setminus \mathcal{R}$.

(3) Line km has power flow $P^{\mathrm{H}}_{km} + jQ^{\mathrm{H}}_{km}$ when H $=$ R *and also when* H $=$ T *and $km \notin \mathcal{L}$.*

(4) When H $=$ R *(reported data) the solution is fully feasible, i.e. it satisfies voltage, generation, phase angle and power flow limits.*

(5) When H $=$ T *(true data) the solution satisfies voltage, generator and phase angle limits, but only satisfies power flow limits on lines km with both $k, m \in \mathcal{A}^C \cup \partial\mathcal{A}$. The solution is also consistent with lines in $\mathcal{L}$ being cut.*

Proof. Property (3) follows from constraints (3.7m)-(3.7p). Hence, (1) and (2) follow from constraints (3.7c)-(3.7f). Properties (4)-(5) follow from constraints (3.7i)-(3.7l). □

As a corollary to (1)-(2) of Lemma 1, a feasible solution to problem (3.7) satisfies, exactly, power-injection consistency, i.e. condition (s.4) above.

Lemma 2. *Consider a feasible solution to problem (3.7). The solution is consistent with a secondary-response adjustment of active power generator amounting to Δ units.*

Proof. Follows from constraint (3.7h). □

Lemma 3. *Consider a feasible solution to problem (3.7). Then* **(a)** *the true and reported voltages agree on $\mathcal{A}^C \cup \partial\mathcal{A}$. Further,* **(b)** *the true and reported currents on a line km are identical if $k, m \in \mathcal{A}^C \cup \partial\mathcal{A}$.*

Proof. **(a)** Follows from constraint (3.7b), and **(b)** is a consequence of **(a)**. □

Corollary 4. *Suppose we compute a feasible solution to problem (3.7) whose objective value is strictly greater than $(S_{uv}^{max})^2$. Then the reported solution amounts to an undetectable attack that hides an overload on line uv.*

3.3.2 Computational Viability

Above we have presented a mathematically correct version of the initial attack problem that would lead to an (initially) undetectable attack, via problem (3.7) which is a nonlinear, nonconvex optimization problem, and thus, in principle, a challenging computational task. Nevertheless this problem is similar to the standard ACOPF or PF problem and (at least) a local optimum should be efficiently computable; this expectation is borne out by our experiments. Strict maximization in (3.7) is *not* required for an attack to be successful (all that is needed is an overload of the line uv), it is easy to see that the objective function (3.7a) can be replaced by a feasibility constraint

$$(P_{uv}^{\mathrm{T}})^2 + (Q_{uv}^{\mathrm{T}})^2 \;>\; (\eta S_{uv}^{\max})^2,$$

for some $\eta > 1$, and solve a feasibility problem instead of an optimization problem.

A broader issue concerns the selection of the sets $\mathcal{A}$ and $\mathcal{L}$. This is a combinatorial problem which is bound to be intractable. In fact [81] describes a number of strong NP-hardness results in the DC setting, e.g. given vectors of phase angles θ and θ' it is NP-hard to compute a set $\mathcal{L}$ such that $B'\theta' = B\theta$ where B' is the bus susceptance matrix of the network with $\mathcal{L}$ removed.

Nevertheless, as discussed in the literature, an attacker may be willing to incur significant computational costs in order to compute a successful attack. While it is reasonable

to assume that an attacker's ability to take physical action or to modify data is limited (see the discussion in [52,64]), not assuming computational intelligence on the part of an attacker amounts to a limitation on the part of the defender.

We separate two distinct issues here: first the identification of the set $\mathcal{A}$, which is done in advance and may be computationally intensive, and second, the solution to problem (3.7) which only requires a few seconds. Let us assume that the attacker has had (undetected) access to system and sensor data long enough to identify a weak sector of the transmission system, i.e. the set $\mathcal{A}$. In this task the attacker would rely on the fact that typical (time- and day-dependent) load and generation profiles for transmission systems are statistically predictable with some accuracy. This fact would help the attacker in the computation of a target set $\mathcal{A}$, perhaps using enumeration, using load estimates in problem (3.7).

Having identified a particular set $\mathcal{A}$, problem (3.7) would be run once again just prior to the attack, now using close estimates of the loads obtained from ambient conditions. Assuming that the attack is perpetrated during a period of slowly changing loads, and not close in time to a generator redispatch[2], the attack will likely be sufficiently numerically accurate so as to become difficult to detect.

3.3.3 Computational Implementation

Our experiments and implementations are based on the MATPOWER test cases. MATPOWER [101] is an open-source package for Matlab that solves optimization problems and power system simulations, such as the DC- and AC-OPF. MATPOWER has a built-in optimization solver called MIPS (MATPOWER Interior Point Solver), a primal/dual interior point method, but other solvers and toolboxes are also supported, such as `fmincon` (from the MATLAB Optimization Toolbox), IPOPT [92], GUROBI [43], CPLEX, MOSEK [6], and others. These solver can be used with MATPOWER if they are properly installed for

[2]A generator redispatch is a request made by the control center to a generator to adjust its real power injection to the network.

running in MATLAB.

In other to implement the type of attacks that have previously been described, we have coded our own optimization/feasibility problems that find solutions to power flow instances. Our code reads test case files that are structured in the same shape as MATPOWER test cases, and creates a MATLAB file where the state variables are the ones described in the previous sections and the objective function and constraints, with their corresponding Jacobian and Hessian, are explicit (in contrast with MATPOWER functions that are difficult to read and modify). The code can be found in the website `github.com/me2533/acopf` and its performance is in some cases better than MATPOWER.

3.3.4 Examples

In the following instances we consider the `case2746wp` and the `case1354pegase` (that have 2746 and 1354 buses, respectively) from the MATPOWER case library [101].

For the first case, the adversary attacks the set of buses

$$\begin{aligned} \mathcal{A} &= \{1137, 1138, 1139, 1141, 1361, 1491\} \\ \text{with} \qquad \mathcal{A}_1 - \partial\mathcal{A}_1 &= \{1137, 1138, 1141, 1491\}. \end{aligned}$$

See Figure 3.1. In this attack the quantity Δ in (3.7h) equals 135.09. We also have $\mathcal{L} = \emptyset$ (no lines are cut). The set of generators participating in secondary response is $\mathcal{R}_1 = \{17,$ 18, 55, 57, 150, 383, 803, 804, 1996$\}$ with participating factors $\alpha_k = 1/9$ for all $k \in \mathcal{R}_1$.

Table 3.1 shows the true and reported flow for lines where the solutions differ, with a strong overload on line $(1361, 1141)$ and $(1138, 1141)$.

Table 3.2 displays the load and generation of the buses involved in the attack.

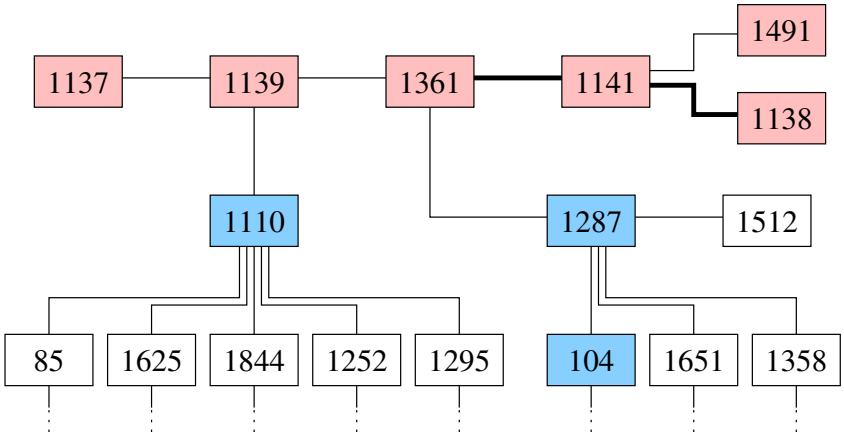

Figure 3.1: Attacked zone $\mathcal{A}_1$ (in red) and its neighborhood. Generators are shown in blue.

Table 3.1: True and reported flow for attacked lines. (Overflow in bold)

bus k	bus m	p^{T}_{km} p^{R}_{km}	q^{T}_{km} q^{R}_{km}	$\|(p^{\mathrm{T}}_{km}, q^{\mathrm{T}}_{km})\|$ $\|(p^{\mathrm{R}}_{km}, q^{\mathrm{R}}_{km})\|$	S^{max}_{km}
1139	1137	3.36 3.36	2.66 2.66	4.29 4.28	114.00
1361	1141	229.01 108.51	10.49 10.49	**229.25** 109.02	114.00
1141	1491	13.46 6.20	2.41 2.39	13.68 6.64	114.00
1141	1138	209.25 98.06	4.44 5.24	**209.29** 98.20	114.00

Table 3.2: Load and generation before and after the attack.

	Before Attack		After Attack			
bus k	$\hat{P}_k^d$	$\hat{Q}_k^d$	$P_k^{d,\mathrm{T}}$	$Q_k^{d,\mathrm{T}}$	$P_k^{d,\mathrm{R}}$	$Q_k^{d,\mathrm{R}}$
1137	0	0	3.36	2.68	14.74	1.37
1138	103.29	29.84	208.91	3.32	20.82	1.46
1139	0	0	3.36	2.68	15.80	1.37
1141	0	0	5.97	2.64	24.20	1.45
1361	0	0	1.58	2.39	91.90	1.45
1491	4.76	1.12	13.45	2.65	20.58	1.48
gen k	$\hat{P}_k^g$	$\hat{Q}_k^g$	P_k^g	Q_k^g	$\alpha_k\Delta$	
17	140.00	120.00	155.01	116.90	15.01	
18	140.00	41.01	155.01	61.33	15.01	
55	130.00	−20.00	145.01	19.29	15.01	
57	130.00	−20.00	145.01	38.33	15.01	
150	90.00	0	105.01	26.65	15.01	
383	21.27	10.54	36.28	13.78	15.01	
803	0	0	15.01	6.47	15.01	
804	0	10.00	15.01	8.27	15.01	
1996	90.00	62.86	105.01	77.28	15.01	

For the second instance we consider the `case1354pegase` and two different attacks. First

$$\mathcal{A}_2 = \{44, 367, 1027, 1172, 1833, 1923, 1973, 2372, 2458, 2644, 2919, 2928, 3243, 3543, 3610, 3657, 3855, 4885, 5308, 5477, 5648, 7148, 7865, 8104, 8722, 9191\}$$

with responding generator set $\mathcal{R}_2 = \{352, 757, 1794, 2421, 2816, 4918, 7267, 7808\}$. See Figure 3.2. The participating factors are the same for all generators.

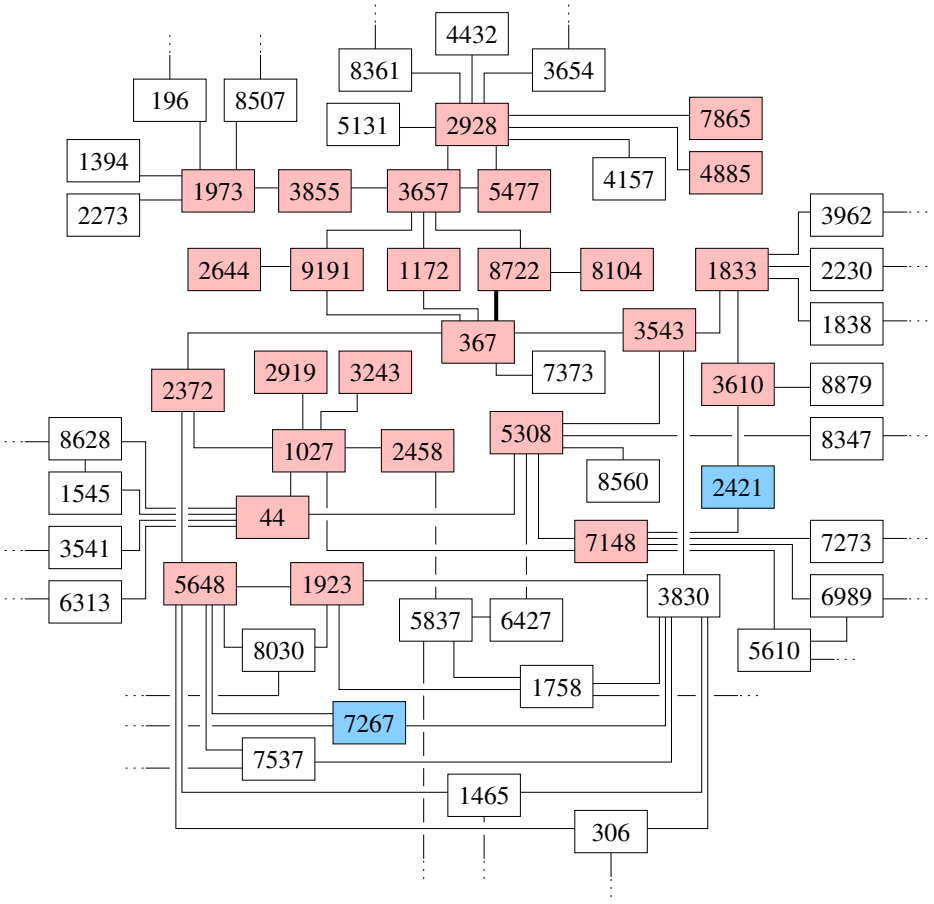

Figure 3.2: Attacked zone $\mathcal{A}_2$ (in red) and its neighborhood. Generators are shown in blue.

Table 3.3 shows the true and reported flow for lines where the solutions differ —with overloads of 53%—, while Table 3.4 displays the load and generation of the buses involved in the attack.

Table 3.3: True and reported flow for attacked lines. (Overflow in bold)

bus k	bus m	p^{T}_{km} p^{R}_{km}	q^{T}_{km} q^{R}_{km}	$\|(p^{\mathrm{T}}_{km}, q^{\mathrm{T}}_{km})\|$ $\|(p^{\mathrm{R}}_{km}, q^{\mathrm{R}}_{km})\|$	S^{max}_{km}
5477	3657	−75.24 −119.01	15.42 18.14	76.80 120.39	491.00
5477	2928	75.24 74.35	−6.50 −8.66	75.52 74.85	433.00
367	1172	384.37 401.00	−106.48 −104.03	398.85 414.28	529.00
367	9191	341.06 404.27	−89.26 −80.36	352.55 412.18	529.00
367	8722	809.96 418.41	−66.76 −81.99	**812.70** 426.37	529.00
3657	3855	146.49 192.83	−48.88 −55.15	154.43 200.56	529.00
3657	3855	162.87 214.05	−48.21 −53.35	169.85 220.59	491.00
3657	2928	108.03 127.77	−27.27 −32.80	111.42 131.92	395.00
3657	2928	112.07 132.54	−31.88 −38.27	116.52 137.95	376.00
3657	1172	−382.70 −298.52	118.57 124.10	400.65 323.29	491.00
3657	9191	−339.80 −271.11	99.37 103.01	354.04 290.02	453.00
3657	8722	117.80 −277.51	99.39 105.57	154.12 296.91	414.00
8104	8722	−0.29 −67.70	0.24 −6.08	0.37 67.97	491.00
2644	9191	0.00 −66.39	0.23 −6.01	0.23 66.67	∞
4885	2928	0.00 −13.49	−0.00 −0.63	0.00 13.50	281.00
3855	1973	195.73 215.31	6.21 −5.72	195.82 215.39	529.00
3855	1973	113.54 124.26	−6.52 −14.44	113.73 125.09	491.00
2928	7865	−0.00 14.67	0.00 −2.73	0.00 14.92	395.00

Table 3.4: Load and generation before and after the attack.

	Before Attack		After Attack			
bus k	$\hat{P}_k^d$	$\hat{Q}_k^d$	$P_k^{d,\mathrm{T}}$	$Q_k^{d,\mathrm{T}}$	$P_k^{d,\mathrm{R}}$	$Q_k^{d,\mathrm{R}}$
367	65.17	−51.50	0	0	311.70	3.89
1027	−5.91	−2.33	0	0	0	0
1172	0	0	0	0	100.80	8.15
1923	192.11	45.50	166.17	0	166.17	0
1973	210.10	66.20	0	49.00	30.27	29.06
2458	132.00	26.00	128.88	0	128.88	0
2644	0	0	0	0	66.39	6.24
2928	0	0	0	14.22	11.10	2.03
3610	185.80	31.80	0	0	0	0
3657	131.00	1.30	0	0	60.91	10.61
3855	−38.73	−15.32	0	20.56	67.17	28.79
4885	166.40	50.40	0	1.49	13.49	2.12
5308	0	0	0	11.08	0	11.08
5477	0	0	0	12.88	44.67	12.32
7148	0	0	0	63.08	0	63.08
7865	84.99	19.60	0	5.23	14.67	2.49
8104	127.20	40.80	0.29	0	67.70	6.32
8722	0	0	924.58	0	71.99	6.16
9191	0	0	0	0	65.55	6.07
gen k	$\hat{P}_k^g$	$\hat{Q}_k^g$	P_k^g	Q_k^g	$\alpha_k\Delta$	
352	870.56	22.55	867.39	−297.68	−3.18	
757	120.00	51.30	116.82	51.30	−3.18	
1794	777.83	−3.27	774.66	−288.66	−3.18	
2421	100.00	−12.45	96.82	37.39	−3.18	
2816	746.19	7.62	743.02	−297.74	−3.18	
4918	80.00	37.27	76.82	−12.57	−3.18	
7267	100.00	23.18	96.82	77.34	−3.18	
7808	797.28	−1.84	794.10	926.35	−3.18	

For the last attack example, also in the `case1354pegase`, we set

$$\begin{aligned}\mathcal{A}_3 \;=\; \{&174, 305, 953, 1035, 1311, 1817, 1965, 2365, 2526, 3579, 3613, 3649, 3697,\\ &3794, 4504, 4874, 5106, 5469, 6555, 6901, 7903, 7905, 8180, 8373, 8748, 8931\}\end{aligned}$$

with responding generator set $\mathcal{R}_3 = \{564,\ 1001,\ 7466\}$. See Figure 3.3. The participating factors are the same for all generators.

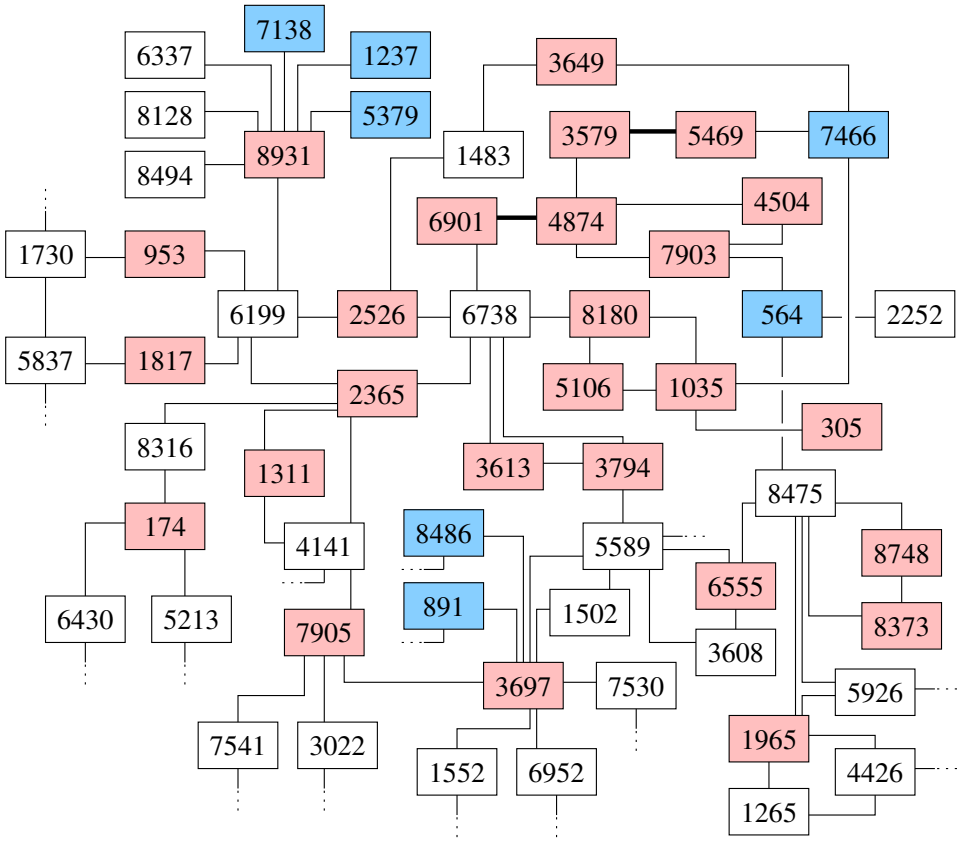

Figure 3.3: Attacked zone $\mathcal{A}_3$ (in red) and its neighborhood. Generators are shown in blue.

Table 3.5 shows the true and reported flow for lines where the solutions differ —with line overloads of 66%—, while Table 3.6 displays the load and generation of the buses involved in the attack.

Table 3.5: True and reported flow for attacked lines. (Overflow in bold)

bus k	bus m	p_{km}^{T} p_{km}^{R}	q_{km}^{T} q_{km}^{R}	$\|(p_{km}^{\mathrm{T}}, q_{km}^{\mathrm{T}})\|$ $\|(p_{km}^{\mathrm{R}}, q_{km}^{\mathrm{R}})\|$	S_{km}^{max}
4874	4504	−54.84 127.62	−17.66 −6.05	57.61 127.76	453.00
4874	3579	−241.10 2.65	21.74 16.17	242.08 16.39	376.00
4874	3579	−254.86 3.09	27.21 17.07	256.31 17.35	∞
4874	7903	−202.23 304.73	−56.32 −37.51	209.92 307.03	338.00
4504	7903	−56.89 18.03	−13.90 −16.39	58.56 24.37	∞
3579	5469	−498.23 −120.65	34.98 21.26	**499.46** 122.51	491.00
6901	4874	953.11 569.95	236.71 88.70	**982.06** 576.81	591.00

Table 3.6: Load and generation before and after the attack.

	Before Attack		After Attack			
bus k	$\hat{P}_k^d$	$\hat{Q}_k^d$	$P_k^{d,\mathrm{T}}$	$Q_k^{d,\mathrm{T}}$	$P_k^{d,\mathrm{R}}$	$Q_k^{d,\mathrm{R}}$
174	370.50	−17.00	0	0	0	0
305	103.00	12.00	0	0	0	0
953	108.70	−12.15	0	0	0	0
1035	382.50	57.90	0	0	0	0
1311	37.94	15.20	0	0	0	0
1817	108.91	33.70	0	0	0	0
1965	248.70	8.10	0	0	0	0
2365	37.94	15.40	0	0	0	0
2526	238.45	46.86	0	0	0	0
3579	277.72	7.35	0	0	126.38	15.21
3613	152.50	14.40	0	0	0	0
3649	0	0	155.13	0	155.13	0
3697	−32.71	−13.76	110.41	40.10	110.41	40.10
3794	165.30	7.40	0	0	0	0
4504	170.10	−9.30	1.83	0	108.55	7.80
4874	−0.58	−53.03	1702.23	0	130.51	12.96
5106	39.94	13.00	0	0	0	0
5469	106.60	−18.80	0	0	385.57	36.48
6555	−37.81	−21.64	0	0	0	0
6901	−37.84	−19.23	0	0	383.16	148.01
7903	0	0	59.13	0	639.63	6.27
7905	347.10	87.20	676.25	0	676.25	0
8180	291.60	47.00	0	0	0	0
8373	216.30	−6.90	0	0	0	0
8748	279.70	4.60	0	0	0	0
8931	−0.56	0	712.85	0	712.85	0
gen k	$\hat{P}_k^g$	$\hat{Q}_k^g$	P_k^g	Q_k^g	$\alpha_k\Delta$	
564	100.00	−19.35	60.04	18.47	−39.96	
1001	120.00	55.02	80.04	55.02	−39.96	
7466	100.00	−59.77	60.04	75.85	−39.96	

3.4 The Follow-Up Phase

Following the initial attack, the attacker needs to dynamically perturb the attack data so as to produce a realistic data stream that is both consistent and continues to hide the overload.

In analogy to our notation for the initial problem, at time $t > 0$ following the attack, we denote by $V_k^{\mathrm{R}}(t)$ and $V_k^{\mathrm{T}}(t)$ be reported and true voltages at t and similarly with currents. Reported data for $\mathcal{A}$ will be manufactured by the attacker aiming to approximately satisfy current-voltage consistency (3.1) and power-injection consistency (3.2).

In addition, in this work we assume that the attack is perpetrated when ambient conditions (in particular loads) are, approximately, constant. Let us denote by $V_k^{\mathrm{R}}(0)$ the voltage at a bus k computed by the initial attack (3.7), i.e.

$$V_k^{\mathrm{R}}(0) \doteq |V_k^{\mathrm{R}}|e^{j\theta_k^{\mathrm{R}}}$$

and likewise define the current $I_{km}^R(0)$ on line km. The statement that ambient conditions are approximately constant, post-attack, can be informally rephrased as

$$V_k^{\mathrm{R}}(t) \approx V_k^{\mathrm{R}}(0) \; \forall\, k, \text{ and } I_{km}^{\mathrm{R}}(t) \approx I_{km}^{\mathrm{R}}(0) \; \forall\, km. \tag{3.8}$$

If ambient conditions are approximately constant, (3.8) will hold (statistically) for any bus k and line km not in the attacked zone $\mathcal{A}$ but are otherwise a requirement for the attacker.

Two types of attack have been used in the literature. First, the "noisy data" attack in our setting works as follows:

Template 2. *Noisy Data Attack*

At time $t > 0$ the attacker reports at each bus $k \in \mathcal{A}$ a voltage

$$V_k^{\mathrm{R}}(t) \;=\; V_k^{\mathrm{R}}(0) + \boldsymbol{\nu_k(t)}.$$

Here $\boldsymbol{\nu_k(t)}$ is a random phasor drawn from small variance, zero mean distribution[a]. Likewise the attacker reports for each line km with both ends in $\mathcal{A}$, currents

$$\begin{pmatrix} I_{km}^{\mathrm{R}}(t) \\ I_{mk}^{\mathrm{R}}(t) \end{pmatrix} \;=\; \begin{pmatrix} I_{km}^{\mathrm{R}}(0) \\ I_{mk}^{\mathrm{R}}(0) \end{pmatrix} + \begin{pmatrix} \boldsymbol{\mu_{km}(t)} \\ \boldsymbol{\mu_{mk}(t)} \end{pmatrix} \tag{3.9}$$

where $\boldsymbol{\mu_{km}(t)}, \boldsymbol{\mu_{mk}(t)}$ are drawn from zero mean distributions with small variance.

[a] We use boldface to indicate random variables.

Note that these definitions satisfy requirement (3.8), and approximately satisfy current-voltage consistency. As a functionally equivalent alternative to (3.9) the attacker could simply set

$$\begin{pmatrix} I_{km}^{\mathrm{R}}(t) \\ I_{mk}^{\mathrm{R}}(t) \end{pmatrix} \;=\; Y_{km} \begin{pmatrix} V_k^{\mathrm{R}}(t) \\ V_m^{\mathrm{R}}(t) \end{pmatrix}, \tag{3.10}$$

our analyses below apply to either form.

A second form of attack that has been considered is the **data replay** attack. Here the attacker supplies a previously observed (or computed) pair of time series $V^{\mathrm{R}}(t)$ and $I^{\mathrm{R}}(t)$ for buses and lines within the set $\mathcal{A}$.

Discussion

The reader may recall that in the initial attack computation we enforced that reported voltages in $\partial\mathcal{A}$ are exact, i.e. equal to the true voltages (constraint (3.7b)). In the time-varying phase this condition is necessarily relaxed by the attacker, though this action carries the risk (to the attacker) that current-voltage consistency will not hold, statistically,

for some line km with $k \in \partial\mathcal{A}$ and $m \notin \mathcal{A}$. Thus e.g. in the noisy data attack template given above the distributions for the $\boldsymbol{\nu_k(t)}$, $\boldsymbol{\mu_{mk}(t)}$ and $\boldsymbol{\mu_{km}(t)}$ should have sufficiently small *variance* relative to the variance of ambient conditions. Further requirements on such variances will be discussed in Section 3.5.3. In any case, when ambient conditions (e.g. loads) are nearly constant, the noisy-data attack may continue to approximately satisfy current-voltage and power-injection consistency and thus remain numerically undetectable. The same holds for the data replay version provided the replayed voltages in $\partial\mathcal{A}$ closely approximate ambient conditions.

In the next section we present defensive mechanisms that dynamically change voltages in a way that is unpredictable by the attacker. The key observation is that a substantial change to voltages in $\partial\mathcal{A}$ will cause the noisy-data attack, applied verbatim as in Template 2, to fail, because of large current-voltage inconsistencies on lines km with $k \in \partial\mathcal{A}$ and $m \notin \mathcal{A}$. Of course, the template need not be applied verbatim, and in particular the attacker may seek to leverage the possibility of sensor error. We will consider this point in the next section.

In [83] current-voltage consistency is used in a different setting: (i) the attacked zone $\mathcal{A}$ is known by the defender, (ii) the attacker only disconnects lines. Under a number of assumptions, in particular that there is a matching between $\mathcal{A}^C$ and $\mathcal{A}$ that covers all buses in $\mathcal{A}$ it is shown that the attack can be accurately recovered.

3.5 Defense

In the above sections we showed that, conceptually at least, it is possible to compute high-fidelity attacks that disguise dangerous network conditions. Other attacks are also potentially conceivable, e.g. impedance changes, transformer tap changes, etc. In this section we describe a generic randomized defense strategy that can deployed when a complex attack is suspected. We will assume that the attack impacts a proper subset $\mathcal{A}$ of the system that is unknown to the control center, as was the case above, though the

generic defense strategy applies under more general attacks as well. The strategy can be summarized by the following template:

Procedure 3. *Random Injection Defense.*

Iterate:

D1: Choose, for each $k \in \mathcal{G}$ a (random) value $\boldsymbol{\delta_k}$ such that $\sum_{g\in\mathcal{G}} \boldsymbol{\delta_k} \approx 0$. Command each generator $k \in \mathcal{G}$ to change its output to $P_k^g + \boldsymbol{\delta_k}$.

D2: Following the generation change in step **D1** identify inconsistencies in the observed sensor readings.

Here, an "inconsistency" is a difference between a voltage value that is reported by the attacker, and a corresponding value that predicted by the defender. We will describe several concrete versions of this idea below. See Procedures 4, 5 and 6.

Each iteration would last several seconds, and statistically significant inconsistencies identified by this scheme would be flagged as potential evidence of an attack. Below we will describe several specific implementations of the random ingredient; randomness is used because the attacker cannot anticipate the random injections and thus will not be able to instantaneously update the sensor readings within $\mathcal{A}$. The above strategy could be AGC-like if only generators $k \in \mathcal{R}$ (the responding generators) are allowed to have $\boldsymbol{\delta_k} \neq 0$ and in general it amounts to a generator redispatch. The strategy in Procedure 3 is likely to succeed, in particular against the noisy data or data replay attacks, if the generation changes result in significant voltage changes across the system. Lemma 5 given below explains why a particular implementation of Procedure 3 attains this goal. An additional point is that an implementation of step **D1** should guarantee safe system operation; this consideration leads to computation of the $\boldsymbol{\delta_k}$ in step **D1** by means of an OPF-like problem.

We note that there is an existing literature on using network resources so as to change power flow physics in order to detect structure or faults. See [15–17, 89, 99]. Indeed, even

though the description of our random defense focuses on power injections, one could also consider other random probing strategies that change power flows, such as adjusting transformer settings, controlled line tripping, and the use of DER (distributed energy resources)[3], storage and FACTS (flexible AC transmission system) devices.

There are several implementations of the generic strategy. Generally the defender wants to make the $|\boldsymbol{\delta_k}|$ large because to first order changes in voltage angles are proportional to $\|\boldsymbol{\delta}\|_2$, and a large change in phase angles is likely to give rise to a significant current-voltage or power-injection inconsistencies in sensor readings in $\partial\mathcal{A}$, as discussed above. This idea forms the basis for a simple, current-consistency based version of Procedure 3 given in Section 3.5.2.

An attacker aware that the random defense strategy is applied may try to replace e.g. the noisy data attack with a more careful manipulation of reported data. For example, the attacker could react to a significant change to voltages in $\partial\mathcal{A}$ by solving a nonlinear, nonconvex system of inequalities designed to guarantee approximate current-voltage and power-injection consistency. In addition, any implied load change within $\mathcal{A}$ must be very small (or it would contradict observed frequencies). Finally the attacker would need to perform this computation very quickly, and repeatedly (because the defense will be applied repeatedly).

This online complex computation could in principle be bypassed by the attacker by considering changes to readings of voltages at buses in $\partial\mathcal{A}$ only; with the remaining voltages in $\mathcal{A}$ computed as in Template 2. We will term this the *enhanced* noisy data attack. We remark that the adversary would still have to maintain AC consistency for lines within $\mathcal{A}$, which is nontrivial. Nevertheless, the ability to adjust readings in $\partial\mathcal{A}$ beyond what is prescribed by the noisy data attack may provide some flexibility for the attacker. However, in Section 3.5.2 we will show that when the random injection defense causes large-enough voltage changes in $\partial\mathcal{A}$, the enhanced noisy-data attack fails. See Lemma 10.

[3]Small local artifacts that generate or store energy, located close to the load point that they serve.

A more practicable alternative (for the attacker) would be to consider arbitrary changes to voltages in buses in $\partial\mathcal{A}$, with the remaining voltages in $\mathcal{A}$ obtained as in Template 2. We will term this the *enhanced* noisy data attack. In Section 3.5.2 we will show that when the random defense causes large-enough voltage changes in $\partial\mathcal{A}$, the enhanced noisy-data attack fails. See Lemma 10.

A more sophisticated defensive idea, given in Section 3.5.3, changes the *stochastics* of power flow data, in particular voltage covariance, and probes the corresponding properties of the reported data.

Our defensive strategies can be easily adjusted if sensors are not available throughout the system. Of course, the fewer the sensors the more limited the impact of the defense. Indeed, some interesting work (using the standard, DC-equation state estimation) precisely seeks to perform system identification post-attack when only limited sensor information is available [80–83]. Note that the attack problem becomes easier (for the attacker) if sensors are not widespread. The attacks computed in Section 3.3.4 do assume sensors at every bus, and yet succeed even in large-scale cases.

3.5.1 Controlling Voltages Through Generation Changes

As discussed above, a goal of the defense is to produce large voltage angle changes in buses in $\mathcal{A}$, with the intention of revealing inconsistencies in reported data on lines between $\partial\mathcal{A}$ and $\mathcal{A}^C$. The defender, of course, does not know the set $\mathcal{A}$ and thus it is of interest to understand when the voltage at any given bus can be changed by appropriately choosing the injections δ.

In this section we address the task of changing voltage angles through injections. First, we will argue by using the DC power flow approximation (2.19), that Procedure 3 does succeed in changing phase angles (see Lemmas 5 and 6). In Section 3.5.1.1 we will present experiments under the AC power model that verify the DC-based results. And in Section 3.5.2 we further argue that the voltage changes are large enough to overcome sensor error.

The defensive strategy that we develop, as a specific implementation of the random injection defense Procedure 3, assumes that there is a known set $\mathcal{T}$ of generator buses that are known to be "trusted", that is to say, we can assume that data from buses in $\mathcal{T}$ is known to be unmodified. This concept is not new; see [25, 33, 40, 52] for related discussions. Without such an assumption the entire suite of signals received by the control centers could be falsified and it is questionable whether any meaningful attack reconstruction can be performed. The following template describes the strategy:

Procedure 4. *Pairs-Driven version of Procedure 3.*

At each execution of step **D1**, select a random pair of generator buses s and t, both in $\mathcal{T}$, as well as random $\boldsymbol{\Gamma} > 0$, and use the injections

$$\boldsymbol{\delta_s} = \boldsymbol{\Gamma},\ \boldsymbol{\delta_t} = \boldsymbol{-\Gamma},\ \text{and}\ \boldsymbol{\delta_k} = 0\ \ \forall k \neq s, t. \tag{3.11}$$

In an application of this defense, define

$$\hat{\boldsymbol{P}}^g \ = \ P^g + \boldsymbol{\delta}$$

We denote by $\hat{B}$ the bus susceptance matrix of the network, after the attack. This matrix will be different from the original bus susceptance matrix B in case of a topology or susceptance attack; thus the control center does not know $\hat{B}$. Recall that as stated above we are assuming that the network remains connected after the attack.

Lemma 5. *Suppose $\hat{B}\theta = P^g - P^d$, and $\hat{B}\hat{\boldsymbol{\theta}} = \hat{\boldsymbol{P}}^{\boldsymbol{g}} - P^d$. Let $k \neq t$ be a bus such that the post-attack network contains a path between s and k that does not include t. Then*

$$\hat{\boldsymbol{\theta}}_{\boldsymbol{k}} - \hat{\boldsymbol{\theta}}_{\boldsymbol{t}} \ > \ \theta_k - \theta_t. \tag{3.12}$$

Proof. Equation (3.12) does not change if we subtract from every $\hat{\boldsymbol{\theta}}_{\boldsymbol{h}}$ any constant, and likewise with the θ_h. Thus, without loss of generality $\hat{\boldsymbol{\theta}}_{\boldsymbol{t}} = \theta_t = 0$. Under this assumption (3.12) reads:

$$\hat{\boldsymbol{\theta}}_{\boldsymbol{k}} - \theta_k > 0. \tag{3.13}$$

Let M be the set of buses $p \neq t$ such that

(1) The network contains a path from s to p that avoids t, and

(2) Subject to (1), $\hat{\boldsymbol{\theta}}_{\boldsymbol{p}} - \theta_p$ is *minimum*.

Aiming for a contradiction, we will assume that

$$\hat{\boldsymbol{\theta}}_{\boldsymbol{p}} - \theta_p \;\leq\; 0 \quad \text{for} \quad p \in M. \tag{3.14}$$

Showing that (3.14) is false yields (3.13). For any line km define the flow value $f_{km} = (\hat{\boldsymbol{\theta}}_{\boldsymbol{k}} - \hat{\boldsymbol{\theta}}_{\boldsymbol{m}} - \theta_k + \theta_m)/x_{km}$. Since $\hat{B}(\hat{\boldsymbol{\theta}} - \theta) = \hat{\boldsymbol{P}}^{\boldsymbol{g}} - P^g$, the flow vector f corresponds (under the DC power flow model) to a power flow with $\boldsymbol{\Gamma}$ units of generation at s, $\boldsymbol{\Gamma}$ units of load at t, and zero generation and load elsewhere. Note that for any line km, $f_{km} > 0$ if and only if

$$\hat{\boldsymbol{\theta}}_{\boldsymbol{k}} - \theta_k > \hat{\boldsymbol{\theta}}_{\boldsymbol{m}} - \theta_m. \tag{3.15}$$

This observation implies

$$\hat{\boldsymbol{\theta}}_{\boldsymbol{s}} - \theta_s > 0. \tag{3.16}$$

[To obtain this fact, decompose the flow vector f into a set of path flows from s to t and telescope (3.15) along any such path.] Pick any $p \in M$ and let P be a path from s to p that avoids t. Say $P = v_0, v_1, \ldots, v_i$ where $v_0 = s$ and $v_i = p$, and let h be smallest such that $v_h \in M$. By (3.16) $s \notin M$, i.e., $h > 0$. Then by definition of h, $\hat{\boldsymbol{\theta}}_{\boldsymbol{v_{h-1}}} - \theta_{v_{h-1}} > \hat{\boldsymbol{\theta}}_{\boldsymbol{v_h}} - \theta_{v_h}$, i.e. $f_{v_{h-1},v_h} > 0$. But by assumption $v_h \neq t$. So there exists some line v_h, m such that $f_{v_h,m} > 0$. Therefore using the assumption $\hat{\boldsymbol{\theta}}_{\boldsymbol{k}} - \theta_k \leq 0$ for all $k \in M$, $v_h \in M$, and (3.15),

$$0 \geq \hat{\boldsymbol{\theta}}_{\boldsymbol{v_h}} - \theta_{v_h} > \hat{\boldsymbol{\theta}}_{\boldsymbol{m}} - \theta_m. \tag{3.17}$$

So $m \neq t$, and as a result by construction there is a path from s to m that avoids t. But then (3.17) contradicts the fact that $v_h \in M$. □

Lemma 6. *Suppose k is any bus and that there are at least two generators available to implement Procedure 3. Then a pair s, t satisfying the assumptions of Lemma 5 exists.*

Proof. Choose $s \in \mathcal{R}$ such that s is closest to k. □

Note: if $\mathbf{\Gamma}$ is chosen negative in (3.11), then instead of (3.12) we obtain $\hat{\boldsymbol{\theta}}_{\boldsymbol{k}} - \hat{\boldsymbol{\theta}}_{\boldsymbol{t}} < \theta_k - \theta_t$ through essentially the same proof.

For future reference, we state the following analogue of Lemma 5, with similar proof (omitted).

Lemma 7. *Let suppose θ and $\hat{\boldsymbol{\theta}}$ be as in Lemma 5. Let $k \neq t$ be a bus such that in the post-attack network every path between s and k must include t. Then*

$$\hat{\boldsymbol{\theta}}_{\boldsymbol{k}} - \hat{\boldsymbol{\theta}}_{\boldsymbol{t}} = \theta_k - \theta_t. \tag{3.18}$$

To analyze the pairs-driven defense we will express phase angles using t as the reference bus; see equation (2.20). Thus $\hat{B}\hat{\boldsymbol{\theta}} = \hat{\boldsymbol{P}}^{\boldsymbol{g}} - P^d$ has a unique solution with $\hat{\boldsymbol{\theta}}_{\boldsymbol{t}} = 0$ of the form

$$\hat{\boldsymbol{\theta}} = \breve{B}_t(\hat{\boldsymbol{P}}^{\boldsymbol{g}} - P^d) = \theta + \breve{B}_t\boldsymbol{\delta} \tag{3.19}$$

where $\breve{B}_t$ is an appropriate pseudo-inverse of $\hat{B}$. (In this expression θ is also written with respect to t). As a result of Lemma 5 we have:

Lemma 8. *Let k be any bus. Then, with high probability we will have $\hat{\boldsymbol{\theta}}_{\boldsymbol{k}} = \theta_k + \beta_k\mathbf{\Gamma}$ for some value $\beta_k > 0$.*

Proof. Let k be any bus. Then with high probability, at multiple iterations of the pairs-driven defense the buses k, s and t will satisfy the conditions of Lemma 5. Let us assume that bus t is the reference bus. Then $\hat{\boldsymbol{\theta}}_{\boldsymbol{t}} = \theta_t = 0$, and thus $\hat{\boldsymbol{\theta}}_{\boldsymbol{k}} > \theta_k$. Thus, (by (3.19)) $\hat{\boldsymbol{\theta}}_{\boldsymbol{k}}$ is as desired. □

This results suggests the following detection paradigm:

Procedure 5. *Pairs-driven detection criterion*

As Procedure 4 iterates, for each bus k, estimate the correlation between $\hat{\boldsymbol{\theta}}_{\boldsymbol{k}}$ and $\boldsymbol{\Gamma}$. The defensive procedure terminates when all these estimates are stable. At that point, any bus k whose correlation coefficient is nonpositive is flagged as suspicious.

Lemma 9. *With high probability the pairs-driven defense will defeat the noisy data and data replay attacks in the sense that each bus whose data is modified will be flagged, and any bus that is not attacked will not be flagged.*

Proof. Let k be any bus. Then with high probability by Lemma 8 the control center expects that $\hat{\boldsymbol{\theta}}_{\boldsymbol{k}} = \theta_k + \beta_k \boldsymbol{\Gamma}$ for some value $\beta_k > 0$ unknown to the control center. This fact yields the desired result since in either attack case, if $k \in \mathcal{A}$ the signals produced by the attacker at bus k will not have the stated form. □

Lemmas 5 through 9 assume the DC power flows model, which is only a first-order approximation to the AC model we consider here. In the next section we perform numerical experiments, under the AC power flows model, of the random injection defense 3. A separate issue concerns ambient *noise*; we need voltage changes to overcome currently found noise levels in measurements. This issue is taken up in Section 3.5.2.

3.5.1.1 Numerical experiments using AC power flows

The above discussion concerns DC power flows. In order to investigate how voltages change under injection changes, under AC power flows, we perform a experiments using examples from the Matpower library [101]. For each system we perform ten experiments. In each experiment we compute an AC power flow which is constrained to satisfying the given voltage bounds at all generator buses, but not at load buses, as well as power injection constraints and generator limits, while allowing large injection changes in a random subset of generators. For a non-generator bus k, let V_k^b be its voltage in the

base case (i.e. the Matpower case), and let $V_{i,k}$ be its voltage in experiment $i = 1, \ldots, 10$. Finally, define

$$\text{score}(k) \doteq \max_{1 \le i \le 10} \frac{|V_{i,k} - V_k^b|}{|V_k^b|}.$$

In Table 3.7, "Min Score" is the minimum score across all non-generator buses. Thus the table provides experimental verification for substantial AC voltage changes under random generator injections.

Table 3.7: AC Voltage Changes

Case	Min Score	Average Score
`case118`	11.61%	32.77%
`case1354pegase`	7.62%	51.00%
`case2746wp`	5.00%	10.09%

3.5.2 Overcoming Sensor Error, and the Current-Voltage Defense

If sensor misestimation (i.e., *error*) is present, a strategy based on Procedure 3 may fail to detect data inconsistencies if the random power injections cause voltage changes that are too small as compared to the error. In order to derive a version of Procedure 3 that deals with this issue, we next describe a particular implementation of step **D2** which relies on the current-voltage consistency condition (3.1) which takes into account the possibility of sensor error. This implementation will take into account the possibility of sensor *error*. Whereas above a phasor (voltage or current quantity) ϕ had a true value ϕ^{T} (the physical value) and a reported value ϕ^{R} (the value received by the control center), now we will have the *sensed* value ϕ^{S} which is the value actually produced by the sensor.

Due to sensor error, sensed and true data may differ. For a phasor ϕ define $\text{err}(\phi) \doteq \phi^{\mathrm{S}} - \phi^{\mathrm{T}}$. In the PMU setting, the TVE (total vector error) criterion [70, 84] guarantees

that

$$|\mathrm{err}(\phi)| \ < \ \tau|\phi^{\mathrm{T}}|, \tag{3.20}$$

where $0 < \tau < 1$ is a tolerance. Standards enforce $\tau = 1\%$, though experimental testing of PMUs shows far smaller errors [36]. From (3.20) we have that

$$\begin{aligned} |\phi^{\mathrm{S}}| - |\phi^{\mathrm{T}}| \ \leq \ |\phi^{\mathrm{S}} - \phi^{\mathrm{T}}| \ < \ \tau|\phi^{\mathrm{T}}| \\ \text{and} \qquad |\phi^{\mathrm{T}}| - |\phi^{\mathrm{S}}| \ \leq \ |\phi^{\mathrm{S}} - \phi^{\mathrm{T}}| \ < \ \tau|\phi^{\mathrm{T}}|, \end{aligned}$$

therefore, we obtain

$$(1-\tau)|\phi^{\mathrm{T}}| \ < \ |\phi^{\mathrm{S}}| \ < \ (1+\tau)|\phi^{\mathrm{T}}| \tag{3.21a}$$

$$|\mathrm{err}(\phi)| \ < \ \frac{\tau}{1-\tau}|\phi^{\mathrm{S}}|. \tag{3.21b}$$

We will describe three sensor-error-aware voltage-current consistency criteria. An important point is that the current-voltage consistency condition (3.1), combined with estimations of possible sensor error, yields a nonlinear relationship, and appropriately reformulation of this relationship can render useful benefits. To simplify notation we will drop the "(t)" from phasors though it should be understood throughout. For a line km write

$$Y_{km} \ = \ \begin{bmatrix} \mathrm{y}_{kk} & \mathrm{y}_{km} \\ \mathrm{y}_{mk} & \mathrm{y}_{mm} \end{bmatrix}.$$

Criterion 1. We have that $I^{\mathrm{T}}_{mk} = \mathrm{y}_{mk}V^{\mathrm{T}}_k + \mathrm{y}_{mm}V^{\mathrm{T}}_m$. Write $\mathrm{z}_{mk} \doteq [\mathrm{y}_{mk}]^{-1}$. Hence

$$V^{\mathrm{S}}_k - \mathrm{z}_{mk}(I^{\mathrm{S}}_{mk} - \mathrm{y}_{mm}V^{\mathrm{S}}_m) \ = \ \mathrm{err}(V_k) \ - \ \mathrm{z}_{mk}(\mathrm{err}(I_{mk}) - \mathrm{y}_{mm}\mathrm{err}(V_m)) \tag{3.22}$$

which yields, using (3.20), (3.21), and the triangle inequality

$$\begin{aligned} |V^{\mathrm{S}}_k - \mathrm{z}_{mk}(I^{\mathrm{S}}_{mk} - \mathrm{y}_{mm}V^{\mathrm{S}}_m)| \ &\leq \ |\mathrm{err}(V_k)| + |\mathrm{z}_{mk}|(|\mathrm{err}(I_{mk})| + |\mathrm{y}_{mm}||\mathrm{err}(V_m)|) \\ &< \ \tau|V^{\mathrm{T}}_k| + \frac{\tau|\mathrm{z}_{mk}|}{1-\tau}(|I^{\mathrm{S}}_{mk}| + |\mathrm{y}_{mm}|\,|V^{\mathrm{S}}_m|) \\ &= \ \tau|\mathrm{z}_{mk}(I^{\mathrm{T}}_{mk} - \mathrm{y}_{mm}V^{\mathrm{T}}_m)| + \frac{\tau|\mathrm{z}_{mk}|}{1-\tau}(|I^{\mathrm{S}}_{mk}| + |\mathrm{y}_{mm}|\,|V^{\mathrm{S}}_m|) \\ &\leq \ \frac{2\tau|\mathrm{z}_{mk}|}{1-\tau}(|I^{\mathrm{S}}_{mk}| + |\mathrm{y}_{mm}|\,|V^{\mathrm{S}}_m|). \end{aligned}$$

In summary, Criterion 1 states that the sensed values $V_k^{\rm S}$, $V_m^{\rm S}$ and $I_{mk}^{\rm S}$ ove lines km must satisfy the following inequality:

$$|V_k^{\rm S} - {\rm z}_{mk}(I_{mk}^{\rm S} - {\rm y}_{mm}V_m^{\rm S})| \;<\; \frac{2\tau|{\rm z}_{mk}|}{1-\tau}(|I_{mk}^{\rm S}| + |{\rm y}_{mm}|\,|V_m^{\rm S}|). \tag{3.23}$$

Under Criterion 1, if, statistically, the reported phasors $V_k^{\rm R}, V_m^{\rm R}, I_{mk}^{\rm R}$ fail to satisfy (3.23) line km is flagged as suspicious. A similar analysis concerns $V_k^{\rm R}, V_m^{\rm R}, I_{km}^{\rm R}$. **Remark:** By construction, if $k, m \notin \mathcal{A}$ then line km will not be flagged.

Criterion 2. Proceeding as above we have

$$\begin{aligned} |I_{km}^{\rm S} - {\rm y}_{kk}V_k^{\rm S} - {\rm y}_{km}V_m^{\rm S}| \;&=\; |{\rm err}(I_{km}) - {\rm y}_{kk}{\rm err}(V_k) - {\rm y}_{km}{\rm err}(V_m)| \\ &<\; \frac{\tau}{1-\tau}(|I_{km}^{\rm S}| + |{\rm y}_{kk}||V_k^{\rm S}| + |{\rm y}_{km}||V_m^{\rm S}|). \end{aligned} \tag{3.24}$$

(and similarly with I_{mk}), and

Criterion 3. When line km is a pure impedance line (no transformer), from (2.7) we have that

$$Y_{km} \;=\; \begin{bmatrix} {\rm y}_{kk} & {\rm y}_{km} \\ {\rm y}_{mk} & {\rm y}_{mm} \end{bmatrix} \;=\; \begin{bmatrix} y_{km} + \dfrac{y_{km}^{\rm sh}}{2} & -y_{km} \\ -y_{km} & y_{km} + \dfrac{y_{km}^{\rm sh}}{2} \end{bmatrix},$$

with branch admittance y_{km} and shunt admittance $y_{km}^{\rm sh}$

$$\begin{aligned} |I_{km}^{\rm S} + I_{mk}^{\rm S}| \;&=\; |(I_{km}^{\rm S} - I_{km}^{\rm T}) + (I_{mk}^{\rm S} - I_{mk}^{\rm T}) + ({\rm y}_{kk} + {\rm y}_{mk})V_k^{\rm T} + ({\rm y}_{km} + {\rm y}_{mm})V_m^{\rm T}| \\ &=\; |{\rm err}(I_{km}) + {\rm err}(I_{mk}) + ({\rm y}_{kk} + {\rm y}_{mk})V_k^{\rm T} + ({\rm y}_{km} + {\rm y}_{mm})V_m^{\rm T}| \\ &=\; |{\rm err}(I_{km}) + {\rm err}(I_{mk}) + \frac{y_{km}^{\rm sh}}{2}(V_k^{\rm T} + V_m^{\rm T})| \\ &<\; \frac{\tau}{1-\tau}(|I_{km}^{\rm S}| + |I_{mk}^{\rm S}|) \;+\; \frac{\tau|y_{km}^{\rm sh}|}{2(1-\tau)}(|V_k^{\rm S}| + |V_m^{\rm S}|). \end{aligned} \tag{3.25}$$

If the reported phasors do not satisfy (3.24) or (3.25) then the line is flagged.

Discussion

Note that a line not attacked will not be flagged, as per the TVE condition. Additional criteria can be developed to handle power-injection consistency. To analyze the effectiveness of these criteria, we turn to the *enhanced* noisy data attack discussed in Section 3.4.

To remind the reader, in this type of attack the voltage readings in $\partial\mathcal{A}$ can be arbitrarily adjusted. While this action may create inconsistencies on lines with just one end in $\partial\mathcal{A}$ the attacker may be able to "hide" such inconsistencies if they are small enough relative to sensor error.

We next show that Criterion 1 alone can suffice to defeat the enhanced noisy data attack (i.e. uncover inconsistencies) when voltage angles are sufficiently changed under our random injection defense.

To understand this point, consider a bus $k \in \partial\mathcal{A}$ such that there is a line km with $m \notin \mathcal{A}$ and also a line ka where $a \in \mathcal{A} - \partial\mathcal{A}$. We study an iteration of the random defense which (to simplify notation) we assume begins at time $t = 0$. Consider line ka first. To avoid having line ak flagged, the attacker t will need to manufacture a time series $V_k^{\mathrm{R}}(t)$, $V_a^{\mathrm{R}}(t)$ and $I_{ak}^{\mathrm{R}}(t)$ that (statistically) satisfy (3.23). But under the noisy data attack, on average $V_a^{\mathrm{R}}(t) = V_a^{\mathrm{R}}(0)$ and $I_{ak}^{\mathrm{R}}(t) = I_{ak}^{\mathrm{R}}(0)$. Hence the attacker needs (on average) that

$$\begin{aligned} \frac{2\tau|\mathrm{z}_{ak}|}{1-\tau}(|I_{ak}^{\mathrm{R}}(0)| + |\mathrm{y}_{aa}|\,|V_a^{\mathrm{R}}(0)|) &> |V_k^{\mathrm{R}}(t) - \mathrm{z}_{ak}(I_{ak}^{\mathrm{R}}(0) - \mathrm{y}_{aa}V_a^{\mathrm{R}}(0))| \\ &= |V_k^{\mathrm{R}}(t) - V_k^{\mathrm{R}}(0)|. \end{aligned}$$

Now consider line km. Since $m \notin \mathcal{A}$, $V_m^{\mathrm{R}}(t) = V_m^{\mathrm{S}}(t)$ and $I_{mk}^{\mathrm{R}}(t) = I_{mk}^{\mathrm{S}}(t)$. Also, denote:

- $V_k^{\mathrm{T}}(*)$ = the true voltage at k at the start of the current iteration of the random defense, i.e. the voltage resulting from the injection changes in step **D1**. Then, assuming unbiased sensor errors and zero-mean ambient noise, $V_m^{\mathrm{T}}(*)$ will equal the expectation of $V_m^{\mathrm{T}}(t)$ during the iteration.

- Likewise define the current $I_{mk}^{\mathrm{T}}(*)$.

Hence the attacker needs (on average) that

$$\begin{aligned} \frac{2\tau|\mathrm{z}_{mk}|}{1-\tau}(|I_{mk}^{\mathrm{T}}(*)| + |\mathrm{y}_{mm}|\,|V_m^{\mathrm{T}}(*)|) &> |V_k^{\mathrm{R}}(t) - \mathrm{z}_{mk}(I_{mk}^{\mathrm{T}}(*) - \mathrm{y}_{mm}V_m^{\mathrm{T}}(*))| \\ &= |V_k^{\mathrm{R}}(t) - V_k^{\mathrm{T}}(*)|. \end{aligned}$$

As a result of these observations we have:

Lemma 10. *Consider buses k, a, m as described above. Suppose that*

$$\begin{aligned} &|V_k^{\mathrm{T}}(*) - V_k^{\mathrm{R}}(0)| \\ &> \frac{2\tau|\mathrm{z}_{ak}|}{1-\tau}(|I_{ak}^{\mathrm{R}}(0)| + |\mathrm{y}_{aa}|\,|V_a^{\mathrm{R}}(0)|) + \frac{2\tau|\mathrm{z}_{mk}|}{1-\tau}(|I_{mk}^{\mathrm{T}}(*)| + |\mathrm{y}_{mm}|\,|V_m^{\mathrm{T}}(*)|) \end{aligned} \tag{3.26}$$

Then it is impossible for the enhanced noisy data attacker to statistically satisfy Criterion 1 on both lines ka and km. □

Comment: This lemma highlights how large changes in voltages caused by the random defense challenge the attacker.

Experiment

Next we describe a set of experiments involving the current-voltage defense applied to the attack given in Section 3.3.4. The current defense was implemented as follows:

- For any generator bus $k \notin \mathcal{R}$, $|\delta_k| \leq \epsilon P_k^g$. We used values $\epsilon = 0.01, 0.05$.
- The set of responding generators, $\mathcal{R}$, was of cardinality 200. For $k \in \mathcal{R}$ $|\delta_k|$ can be arbitrarily large. We chose $\delta_k > 0$ with probability 1/2.
- No generator may exceed its limits (voltage or generation), but subject to all these conditions we maximize $\sum_{k\in\mathcal{G}} |\delta_k|$.

In Table 3.8, we perform the above analysis on the lines ($k = 1139, a = 1137$) and ($k = 1139, m = 1110$) with $\tau = 0.01$. "Ratio" is the ratio of the left-hand side to the right-hand side of expression (3.26). We see that the condition for Lemma 10 is amply satisfied. A similar analysis pertains to line $(1141, 1361)$, the other line connecting $\mathcal{A}$ to its complement.

Table 3.8: Current-voltage defense.

	Experiment 1	Experiment 2
ϵ	0.01	0.05
$\sum_{k\in\mathcal{G}} \delta_k^+$	289.01	964.77
$\sum_{k\in\mathcal{G}} \delta_k^-$	174.47	256.04
Line $(k = 1139, a = 1137)$		
$\lvert V_a^{\mathrm{R}}(0)\rvert \angle \theta_a^{\mathrm{R}}(0)$	$1.0919\angle -6.993°$	$1.0919\angle -6.993°$
$I_{ak}^{\mathrm{R}}(0)$	$-0.0275 + 0.0281j$	$-0.0275 + 0.0281j$
Line $(k = 1139, m = 1110)$		
$\lvert V_m^{\mathrm{T}}(*)\rvert \angle \theta_m^{\mathrm{T}}(*)$	$1.0309\angle -7.822°$	$1.0391\angle -7.848°$
$I_{mk}^{\mathrm{T}}(*)$	$0.0905 - 0.4976j$	$0.1289 - 0.4901j$
Voltages at $k = 1139$		
$\lvert V_k^{\mathrm{R}}(0)\rvert \angle \theta_k^{\mathrm{R}}(0)$	$1.0919\angle -6.991°$	$1.0919\angle -6.991°$
$\lvert V_k^{\mathrm{T}}(*)\rvert \angle \theta_k^{\mathrm{T}}(*)$	$1.0104\angle -7.822°$	$1.0187\angle -7.936°$
Lemma 10 applied to bus $k = 1139$		
Ratio	1.913	1.732

3.5.3 Covariance Defense

In this section we describe an elaboration of the random-pairs defense Procedure 4; the elaboration is motivated by the fact that real-world PMU data streams exhibit non-generic stochastic structure in (for example) voltage angles [19, 21, 93]. In particular, covariance matrices across several time scales have very low rank (typically smaller than 10). Our defense will defeat both the noisy-data and data-replay attacks, under appropriate assumptions.

As before, we assume that the buses in a certain set $\mathcal{T}$ are *trusted*. The emphasis of the methods in this section is that we aim to modify the *covariance* matrix of phase angles, whereas the random injection defense in Procedures 3 or 4 change the *average* voltage values. Such a change should prove more difficult for the attacker to correctly counteract since such a correction involves an estimation that requires time, during which

the attacker will be producing incorrect data.

We additionally assume that the attacker's data stochastics are *stationary* (i.e. the parameters of the stochastic process do not change as a function of time). This implies in particular that the attacker does not react to the covariance defense by changing the stochastics. Below we will discuss, however, why reacting to the defense would prove very difficult. We also assume that ambient conditions are also stationary.

In order to describe the defense we need some definitions. For a pair of buses $s, t \in \mathcal{T}$, define the vectors $u^{s,t}$ and $v^{s,t}$ by

$$u_k^{s,t} \doteq \begin{cases} 1, & \text{if } k = s, \\ -1, & \text{if } k = t, \\ 0, & \text{otherwise} \end{cases} \tag{3.27a}$$

$$v^{s,t} \doteq \breve{B}_t u^{s,t}. \tag{3.27b}$$

Formally, the covariance defense works as follows. Let t_1, t_2 be two fixed trusted buses, and let $\mathcal{P}$ be a real-valued probability distribution, with zero-mean and variance $\sigma_{\mathbf{\Gamma}}^2 > 0$. The defense has two phases.

(I) During an initial phase, after suspecting the attack, for $i = 1, 2$, we compute the matrix

$\sigma^2_{\hat{\boldsymbol{\theta}}^{\mathbf{R},i}} \doteq$ covariance matrix of observed phase angles, with respect to reference bus t_i.

(II) After the initial phase, we perform iterations as in Procedure 4, as follows. We randomly (uniformly) choose one pair of buses of the form (s, t) where $s \in \mathcal{T}$ and $t = t_1$ or t_2, and we draw a random value $\mathbf{\Gamma}$ from the distribution $\mathcal{P}$, independently from the stochastics of the attacker and ambient stochastics. We then apply step **D1** of Procedure 4 using this triple $(s, t, \mathbf{\Gamma})$.

Throughout the second phase, we compute the following matrix, for $i = 1, 2$,

- $\sigma^2_{\hat{\boldsymbol{\theta}}^{\mathbf{R},i}} \doteq$ covariance matrix of observed phase angles, with respect to reference bus t_i.

The defense concludes when the estimates for these two matrices become stable.

Procedure 6. *Covariance-driven detection criterion*

At termination of the defense, we flag a bus k as *suspicious* if, for both $i = 1$ and 2, the difference between the (k, k) entry of $\sigma^2_{\hat{\boldsymbol{\theta}}^{\mathbf{R},i}}$ and the corresponding entry of $\sigma^2_{\boldsymbol{\theta}^{\mathbf{R},i}}$ is smaller than λ, defined by

$$\lambda \doteq \frac{\sigma^2_{\mathbf{\Gamma}}}{|\mathcal{T}| - 1}\omega, \text{ where} \tag{3.28a}$$

$$\omega \doteq \min_{s,t,j}\{(v^{s,t}_j)^2 : v^{s,t}_j \neq 0\}. \tag{3.28b}$$

This concludes the description of the defense, with analysis given in Lemmas 11, 13 and 9.

In preparation for those results, suppose that at some point in phase **(II)** the pair (s, t_i) has been selected. Let

(a.1) $\hat{\boldsymbol{\theta}}^{\mathbf{T},i}$ be the vector of true voltage phase angles, using t_i as the reference bus.

(a.2) $\boldsymbol{\theta}^{\mathbf{T},i}$ describe the true vector of voltage phase angles, had the power injections in the defense *not* been applied at that point of time. It is also given using t_i as the reference bus.

Lemma 11 given next concerns the relationship between these last two vectors.

Lemma 11. *Suppose that at some point in* **(II)** *the pair* (s, t_i) *is being used. Then,*

$$\hat{\boldsymbol{\theta}}^{\mathbf{T},i} = \boldsymbol{\theta}^{\mathbf{T},i} + \mathbf{\Gamma} v^{s,t_i}.$$

Proof. Given that the pair (s, t_i) is being used, the injections in the random defense, per equation (3.11) are given by $\boldsymbol{\delta} = \mathbf{\Gamma} u^{s,t_i}$. The result follows from equation (3.19). □

The next result presents a key feature of the covariance of phase angles. Recall that $\mathbf{\Gamma}$ is drawn independent of all other stochastics.

Lemma 12. *For* $i = 1, 2$,

$$\sigma^2_{\hat{\boldsymbol{\theta}}^{\mathbf{T},i}} = \sigma^2_{\boldsymbol{\theta}^{\mathbf{T},i}} + \frac{\sigma^2_{\mathbf{\Gamma}}}{|\mathcal{T}| - 1}\sum_{s \in \mathcal{T} - t_i} v^{s,t_i}(v^{s,t_i})^\top. \tag{3.29}$$

Proof. We proceed by conditioning on the pair (s, t_i) being selected by the defense. Subject to this conditioning, by Lemma 11 the covariance of $\hat{\boldsymbol{\theta}}^{\mathbf{T},\boldsymbol{i}}$ equals

$$\sigma^2_{\boldsymbol{\theta}^{\mathbf{T},\boldsymbol{i}}} + \sigma^2_{\boldsymbol{\Gamma}} v^{s,t_i} (v^{s,t_i})^\top + \operatorname{covar}(\boldsymbol{\theta}^{\mathbf{T},\boldsymbol{i}}, \boldsymbol{\Gamma} v^{s,t_i}).$$

The last term in this expression is zero, by the independence assumption on $\boldsymbol{\Gamma}$. The result follows since each pair is chosen with probability $(|\mathcal{T}| - 1)^{-1}$. □

Lemma 13. *Let k be any bus. Then, for at least one of $i = 1$ or 2, the (k, k) entry of $\sigma^2_{\hat{\boldsymbol{\theta}}^{\mathbf{T},\boldsymbol{i}}}$ is at least as large as the corresponding entry of $\sigma^2_{\boldsymbol{\theta}^{\mathbf{T},\boldsymbol{i}}}$, plus λ —defined as in (3.28a).*

Proof. Without loss of generality, there is a path between k and t_1 that avoids t_2. Note that the pair (s, t_2) with $s = t_1$ is one of the pairs available for the defense. By Lemma 5, we have that $v_k^{t_1,t_2} > 0$ and so $v_k^{t_1,t_2} \geq \omega^{1/2}$. Considering (3.29) for $i = 2$ we see that one of the terms in the sum corresponds to $s = t_1$. As just argued, the (k, k) entry of this term is at least ω. The (k, k) entries in the remaining terms of the sum are nonnegative (since each term is a positive-semidefinite matrix). Thus the result follows. □

Lemma 14. *The suspicious labels computed by the covariance defense are correct.*

Proof. Consider first a bus k that is not attacked. For such a bus, by definition, $\hat{\boldsymbol{\theta}}_k^{\mathbf{T},\boldsymbol{i}} = \hat{\boldsymbol{\theta}}_k^{\mathbf{R},\boldsymbol{i}}$, for both $i = 1, 2$. Thus, by Lemma 13, bus k is not flagged as suspicious. On the other hand, suppose k is attacked. Then under either the noisy-data or data-replay attacks the (k, k) entry of $\sigma^2_{\hat{\boldsymbol{\theta}}^{\mathbf{R},\boldsymbol{i}}}$ will be equal to the corresponding entry of $\sigma^2_{\boldsymbol{\theta}^{\mathbf{R},\boldsymbol{i}}}$, by the stationarity assumption (the attacker does not change stochastics when the defense is implemented). Hence bus k is flagged. □

Remarks:

(1) Recall (3.28a) and (3.28b). The quantity ω depends on the bus susceptance matrix B, only. Hence by choosing $\sigma^2_{\boldsymbol{\Gamma}}$ large enough we can make λ large. Also note that in the above proofs, we can restrict the set $\mathcal{T}$ to a subset of size 2, again helping λ attain large values.

(2) Given a pair (s, t_i) used in the defense, by Lemma 5 any entry v_k^{s,t_i} is positive if there is a path from bus k to s that avoids t_i. Let A denote the set of such buses. By Lemma 7, for $k \notin A$, $v_k^{s,t_i} = 0$.

Thus, in the term $v^{s,t_i}(v^{s,t_i})^\top$ in (3.29) the entire submatrix with rows and columns in A is positive, and the remaining entries in $v^{s,t_i}(v^{s,t_i})^\top$ are zero. By adjusting the proof of Lemma 13 we conclude that for k and m in A, the entry (k, m) of $\sigma^2_{\hat{\boldsymbol{\theta}}^{\mathbf{T}},i}$ is at least as large as the corresponding entry of $\sigma^2_{\boldsymbol{\theta}^{\mathbf{T}},i}$, plus λ. Thus a submatrix of the covariance matrix will change via the defense (and not just the diagonal entries). If the network is guaranteed to be 2-connected [5] one can in fact prove that the entire matrix must change.

The covariance defense has an additional important feature, namely that the sum on the right-hand side of (3.29) has rank $|\mathcal{T}| - 1$ (as shown next) whereas we expect the left-hand side of (3.29) to have low rank.

Lemma 15. *For $i = 1, 2$ the vectors $v^{s,t_i} = \breve{B}_{t_i} u^{s,t_i}$ as in (3.27b) are linearly independent. Hence the second term in (3.29) has rank at least $|\mathcal{T}| - 1$.*

Proof. The $|\mathcal{T}| - 1$ vectors u^{s,t_i} arising from all pairs (s, t_i) under consideration are linearly independent, by construction in (3.27a). Hence, the corresponding vectors v^{s,t_i} are also linearly independent. □

Lemma 15 highlights the challenges faced by the attacker, even if the attacker is aware that the covariance defense is being deployed. The attacker will have to alter reported data in a way consistent with an appropriate rank change, but the attacker does not know the pairs (s_i, t) being used (or the distribution $\mathcal{P}$). Such "learning" would require data observations, i.e. time, during which the attacker is still expected to produce data readings, producing an error trail.

Chapter 4

Learning from PMU data

This chapter describes a statistical study that we have performed using a dataset of PMU measurements that we have obtained from an industrial partner. We detail the nature of the data together with the statistical tools that we have used. Plots and videos help visualize properties of the sampling.

This study was mainly performed under the supervision of Prof. Michael Chertkov while I was visiting Los Alamos National Lab, NM, in September 2018. Dr. Jonatan Ostrometzky has helped to improve the Fourier filters that we have applied.

4.1 Introduction

Data-driven techniques in power systems have at least fifty years of history, starting with static state estimations developed by Schweppe and co-authors [77–79], then transitioning to dynamic state estimation analysis and applications, see e.g. [45, 62, 88] and references therein, and most recently discussed under the umbrella of "big data" as the most significant enabler of power system operations, security and resiliency in the future [9, 49]. (See also related discussion in the description of the US Department of Energy new funding opportunity to "explore the use of big data, artificial intelligence, and machine learning technology to leverage the power of grid sensors" [4].) Many specific questions and

approaches, including but not limited to detection of modes of oscillation and related analysis of stability [1–3, 88], dimension reduction for faster processing and analysis [31], early event detection [94], missing data recovery [38], identification of cyber attacks [37] and real time (online) event detection [57] are among the most recent research thrusts.

On the methodology side data-driven methods developed in other engineering disciplines have been adopted, modified and used for many (e.g. aforementioned) power system applications. Principal component analysis [31, 37, 38, 57, 94], auto-correlation analysis of memory effects [74], and linear model driven spectral analysis of the dynamic state matrix [1–3, 27, 62, 65, 75, 88] are arguably the most popular data-driven techniques currently in use in the power system research.

Even though the sophistication level of the methods already used in power system applications is impressive, coherence and understanding of the potential of new generation of big data methods, driven during the last decade largely through heavy investment of IT industry, is still lacking [4]. We anticipate that many of the most modern methods, especially Deep Learning (DL) and related techniques linked to Machine Learning (ML) and Artificial Intelligence (AI), revolutionary advances in data science and more generally theoretical engineering [8, 12, 42, 56, 76], will impact the power-system operation-room reality in an even more significant ways. However, one problem with applications of the novel methods of DL and alike in sciences and engineering is that they are application agnostic/generic – very effective for many business cases, but lacking "explainability", i.e. intuitive physical/engineering explanations. This significant handicap of the most advanced and recent ML & AI methods slows down development of related applications in power systems. Indeed, power system practitioners would generally not consider as practical any new methods lacking "power systems informed" explanations.

This manuscript takes a step towards closing the gap between the rich variety of methods already developed and utilized in power systems and the yet to be unleashed power of the upcoming Big Data revolution. Specifically, we start walking towards exciting sophistication of DL slowly, from the well-established and intuitive trenches of practical system

engineering. We develop in this manuscript a pragmatic "phasor-detective" approach to analysis of the streaming Phasor Measurement Unit (PMU) data which allows to extract and interpret spatial and temporal correlations in a computationally light fashion and without making any constraining assumptions about origin of the correlations.

Our Contribution

We analyze synchronized historical PMU data recorded at ≈ 200 most significant locations of a US Independent System Operator (ISO) over the course of two years. At each PMU location the data includes complex current and voltage recorded with a millisecond resolution. Given geo-spatial locations of the PMU, but no information about the grid characteristics and layout, we pose the following principal questions: What can we possibly reconstruct from the data stream about the system current ambient behavior? To answer the question we utilize available statistical tools. In relation to preliminary data processing we apply to the raw signal three filtering techniques: moving average, sliding time horizon and Fourier analysis pre-processing. This allows to provide robust identification of the "quiet" periods and also prepare data for subsequent statistical analysis by means of Principal Component Analysis (PCA) and Auto-Correlation Analysis (ACA). We show that the two complementary tools, applied to the raw as well as to pre-processed signal, allow to separate scales and also provide compressed, thus easy to visualize, descriptors for online tracking of current state of the grid in a much broader way than what the current Energy Management System (EMS) actually uses. PCA provides a robust set of indicators which record slow/adiabatic changes on the scale of seconds to ten of seconds and slower. We have observed that only a very few principal modes are significant at any moment of time, even though these modes may be different for voltage amplitude, phase difference and frequency (the three main characteristics) we track. The results do not change when PCA is applied to the filtered signal, consistently with the fact that PCA averages over time but does not catch different-time correlations. ACA is the tool used to analyze the latter, in particular identifying significant, persistent corre-

lations, missed by PCA, at shorter time scales - subseconds-to-seconds. Following ACA curves at different spatial locations we were able to identify nodes where correlations do not decay with time showing significant memory-effects. Remarkably, these nodes with significant memory cluster geographically. We observe two areas in the grid which show especially strong sustainable temporal correlations. We then proceed with ACA analysis of the Fourier-filtered signal. This helps us to identify and localize different harmonics. In particular, we observe (for a particular quiet period) emergence of significant oscillations in the 4-6 Hz range at a small number of nodes. Interestingly, nodes with significant sustainable oscillations are either wind farms, big aggregated loads or mid-size generators. We conjecture that the sustainable oscillations are indicators of malfunction at these critical elements of the grid. We also observe that sustainable oscillations, seen clearly through emergence of a residue in the ACA analysis of the raw signal, disappear when applied to the Fourier-filtered signal (cutting off the 4-6 Hz oscillations). Finally, analyzing spatial cross-correlations (of the residue) we were able to identify a group of nodes with significant inter-dependency.

4.2 Logic and Main Steps

In this part of the work we study the data streams reported by over 200 PMUs operated by an ISO and spanning a period over one year long. Figure 4.1 displays a rendition of the locations of the PMUs using anonymized coordinates. As mentioned in Section 2.4, the data stream from each PMU includes frequency, (complex) current and voltage, reported 30 times per second. Using this data one can obtain real-time estimates of complex power at each location. Working with a data set this large (on the order of 28 TB) presents some obvious challenges; additionally there are specific artifacts that can arise in the data. For example, not all PMUs are always reporting, and occasionally some PMUs exhibit what appears to be errant behavior.

Our work has centered on performing statistical analysis aimed at inferring “struc-

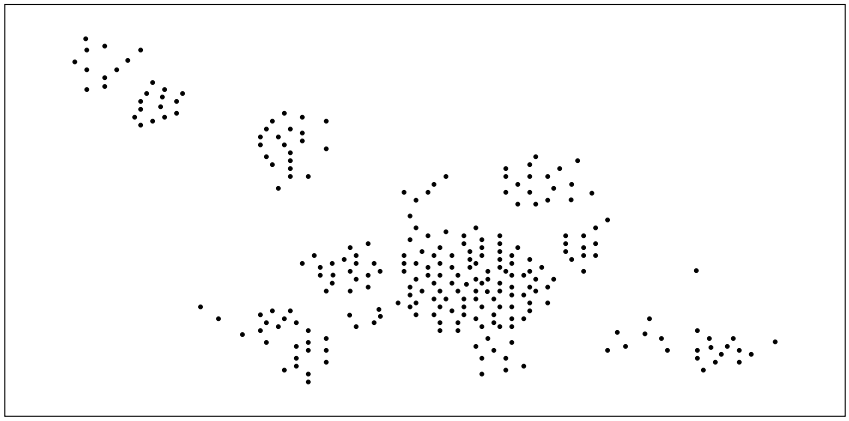

Figure 4.1: Geographical location of PMU's (anonymized coordinates), each mark represents the position of a sensor.

ture" in the underlying transmission system as well as identifying complex behaviors, such as resonance and oscillations. In this manuscript we focus specifically on identification and characterization of "quiet" periods, also known as ambient conditions periods. Oversimplifying (see related discussion below) we consider a period *quiet* if fluctuations around the mean (e.g. characterized in terms of the standard deviations) are smaller than a reasonable pre-defined threshold. This focus on the quiet/ambient periods is motivated by the following considerations:

- The development of a strong understanding of quiet periods and (in particular) efficient online algorithms for recognition of such periods is a necessary step prior to studying less-quiet or even anomalous regimes, for otherwise we risk significant misinterpretation, i.e. errors in online detection of anomalies.

- As will be seen in this paper, the quiet regimes display informative patterns and correlations, all (slowly) time-evolving. Identifying such features is important with regards to:

  - Developing fast and reliable identification techniques.

  - Uncovering hidden malfunction of assets thus providing significant contribu-

tion towards forecasting most probable (and destined to occur) failures.

- The richness of correlations observed in the quiet regime, in fact, suggests that separation of what is normal/quiet from what is anomalous/atypical will be challenging. Even though we observe that quiet regimes dominate, relatively abrupt jumps of moderate size (i.e. jumps exceeding tracked standard deviation by factor of two or three) are rather frequent although overall they account for a relatively short fraction of the stream. As a result, it is rather difficult to find sufficiently long entirely quiet periods in the available data.

- Clearly, understanding quiet vs. volatile behavior will be helpful toward building predictive models for better optimization, control and planning.

- "Cyber-physical" attacks on power systems —like the ones described in Section 3— are a venue where fast and effective learning of (changing) stochastics may prove useful in identifying attacks.

The methodology adopted in this manuscript to identify the quiet periods is explained in Section 4.3.3.

4.3 Description and Averaging of the Time Series

The available data encompasses the period from January 1st 2013 to March 21st 2014. Each of the $N = 240$ PMUs records the following measurements 30 times per second: time of the measurement (GPS tagged), bus ID, voltage amplitude, voltage phase angle, current magnitude, current angle, and frequency. Additionally, the 2-dimensional coordinates of the PMU locations is also available, together with their corresponding nominal voltages. We note that PMUs *report* 30 times/second, but they *sample* at a far higher rate and perform filtering (e.g. anti-aliasing[1]) before reporting.

[1] Aliasing is an effect that causes the measured/sampled signal to be different than the real one. This might be caused for example by precision inaccuracy of the measurement instrument.

We denote a generic scalar or complex measurement (e.g. complex voltage) at PMU location k, at time t by $m_k(t)$. The parameter t will be used to refer to the discrete time sequence with each temporal data point separated from preceding one by the same duration Δ ($1/30^{th}$ of a second in our case).

Typical pre-processing steps in the statistical analysis of data (especially with the goal of analyzing correlations) involve modifications through de-trending, offsetting (subtraction) of moving average, and normalizations. We will apply such techniques below. Specific details are provided next.

4.3.1 Moving Average and Covariance

The *moving average* $\mu^{(m)}$ and *moving variance* $\sigma^{(m)}$ vectors of a sampled series m is computed in the following way,

$$\forall k \in \{1, \ldots, N\},\ \forall t \geq 1 :$$

$$\mu_k^{(m)}(t;\alpha) = \alpha \cdot m_k(t) + (1-\alpha) \cdot \mu_k^{(m)}(t-1;\alpha), \tag{4.1}$$

$$\sigma_k^{(m)}(t;\alpha) = \alpha \cdot \left| m_k(t) - \mu_k^{(m)}(t-1;\alpha) \right|^2 + (1-\alpha) \cdot \sigma_k^{(m)}(t-1;\alpha), \tag{4.2}$$

with some initial values, say $\mu_k^{(m)}(0;\alpha) = 0$ and $\sigma_k^{(m)}(0;\alpha) = 1$. The parameter $\alpha \in (0,1)$ represents the degree of weighting decrease. Note that in (4.2) we are using the moving average defined in (4.1). One can likewise define moving covariance parameters.

Figures 4.2, 4.3 and 4.4 show the sampling of the frequency, voltage angle and voltage magnitude (respectively) for a particular PMU during one minute; together with the corresponding moving average and moving standard deviation for different values of α.

We introduce the zero mean and normalized zero mean data streams

$$\bar{m}_k^{(m)}(t;\alpha) = m_k(t) - \mu_k^{(m)}(t-1;\alpha), \tag{4.3}$$

$$\hat{m}_k^{(m)}(t;\alpha) = \frac{m_k(t) - \mu_k^{(m)}(t-1;\alpha)}{\sqrt{\sigma_k^{(m)}(t-1;\alpha)}}, \tag{4.4}$$

obtained from the input stream by making use of the moving average and variance. The zero mean parameters will help us identify quiet periods, as discussed next.

4.3.2 Averaging over Sliding Time Horizon

Consider a number S of measurements, this value corresponds to the memory budget. We define N dimensional vectors of means and variances (averaged over the (last) sliding time horizon of duration S) as follows:

$$\forall k \in \{1,\dots,N\},\ \forall t, \quad \mu_k^{(s)}(t;S) = \frac{1}{S}\sum_{\tau=t-S+1}^{t} m_k(\tau), \tag{4.5}$$

$$\sigma_k^{(s)}(t;S) = \frac{1}{S}\sum_{\tau=t-S+1}^{t} \left| m_k(\tau) - \mu_k^{(s)}(t;S) \right|^2. \tag{4.6}$$

Figures 4.5, 4.6 and 4.7 show the sampling of the frequency, voltage phase angle and voltage magnitude (respectively) for a particular PMU during one minute; together with the corresponding average and standard deviation over sliding time horizon for different values of S.

As before, we also define the re-scaled zero-mean data vectors:

$$\bar{m}_k^{(s)}(t;S) = m_k(t) - \mu_k^{(s)}(t-1;S), \tag{4.7}$$

$$\hat{m}_k^{(s)}(t;S) = \frac{m_k(t) - \mu_k^{(s)}(t-1;S)}{\sqrt{\sigma_k^{(s)}(t-1;S)}}. \tag{4.8}$$

Numerically, we found that a reasonable choice for S, that allows to separate power electronics (milliseconds) and electro-mechanical (seconds) time scales, is the number of readings in 1 second, that is $S = 30$. Similar normalizations are obtained when $\alpha = 0.05$.

4.3.3 Quiet Periods

Given a reference time t and a length parameter Q consider the $N \times Q$ matrix of normalized measurements

$$M(t;\alpha;Q) = [\ \hat{m}_k^{(m)}(\tau;\alpha)\ |\ \forall k \in \{1,\dots,N\},\ \tau \in \{t-Q+1,\dots,t\}\ ],$$

corresponding to the last Q measurements before time t for all buses. We define the period $(t-Q,t]$ as *quiet* if the absolute value of all entries of the matrix $M(t;\alpha;Q)$ is

below some preset threshold. The reason behind this definition is that sudden jumps in the data appear as large values in the normalized time series, whereas normalized values close to zero mean that the data is behaving in a steady way. Effectively, a quiet period is an interval of time where all sensor-reported data behave in a stationary way. Moreover it is relatively cheap to compute $M(t;\alpha;Q)$ from the stream data in an online fashion, as we just have to complete a column at every time that samples are received and keep track and update of the moving variance of each sensor.

Again, in our analysis, we used $\alpha = 0.05$ as the length of the sliding time horizon, and we consider quiet periods spanning 15 minutes. Over a selection of five different days across the database, we compute we compute the matrix $M(t;\alpha;Q)$ and record its maximum absolute value when t spans over the complete day. Just in few cases the maximum was below 10 units, we selected 2-3 intervals between these recorded cases.

Table 4.1 displays the specific time intervals selected as quiet periods that we have analyzed.

Table 4.1: Selected quiet periods.

#	Date	Time Window
1	January 15, 2013	12:13 – 12:28 AM
2	March 10, 2013	12:09 – 12:24 AM
3	March 10, 2013	4:29 – 4:44 AM
4	March 10, 2013	2:46 – 3:01 PM
5	April 3, 2013	1:52 – 2:07 AM
6	April 3, 2013	7:19 – 7:34 AM
7	April 3, 2013	7:26 – 7:41 PM
8	July 30, 2013	1:26 – 1:41 AM
9	July 30, 2013	4:39 – 4:54 PM
10	July 30, 2013	9:24 – 9:39 PM

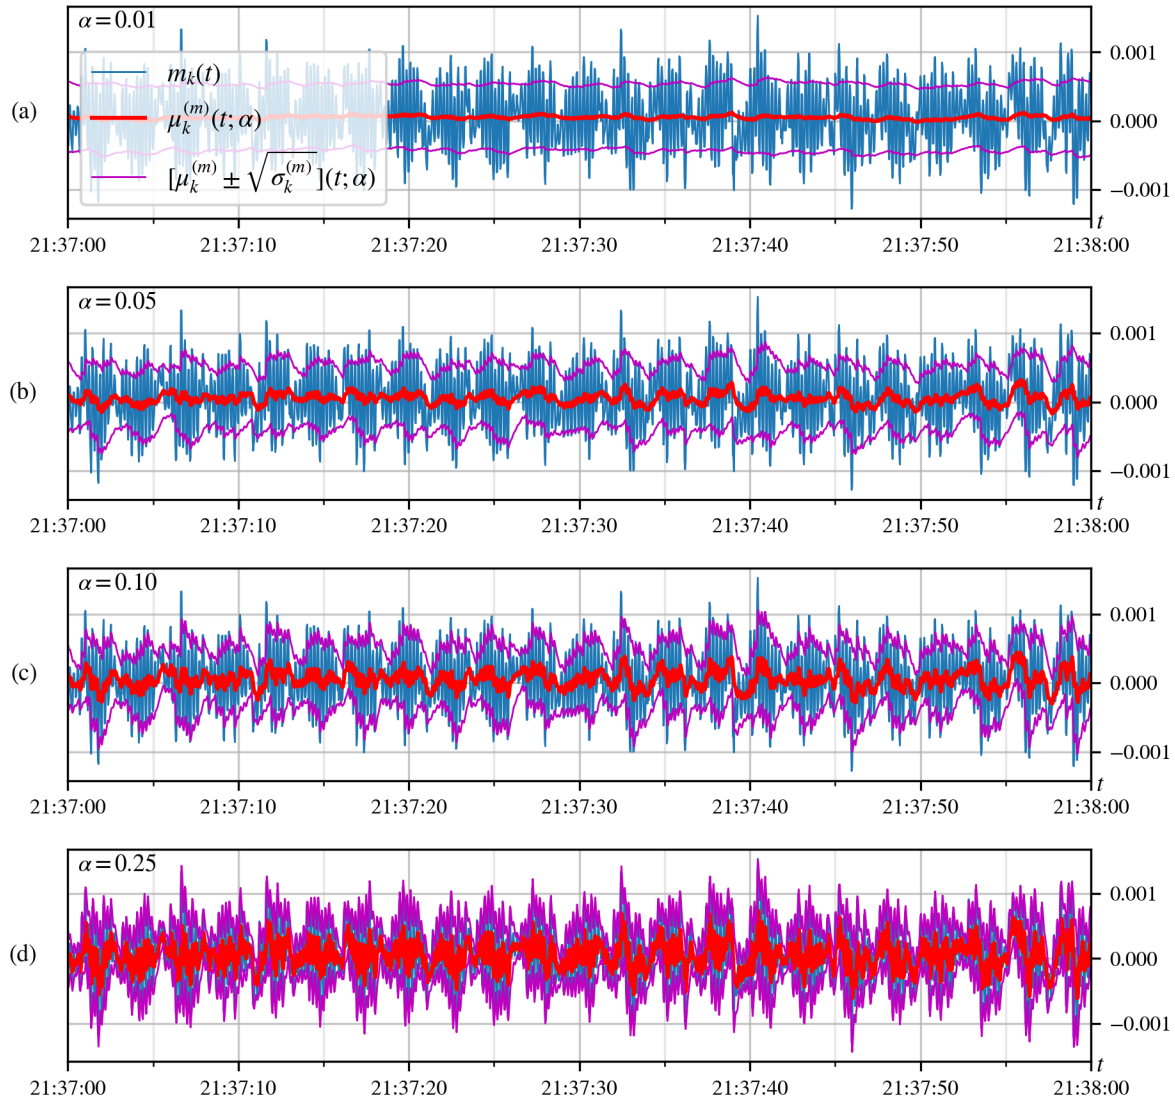

Figure 4.2: Frequency measurements with respect to a reference bus (in blue) of PMU $k = 156$ during one minute, between 9:37 and 9:38 PM on July 30, 2013. Moving average (in red) and a band of one standard deviation (in purple) are also shown for (a) $\alpha = 0.01$, (b) $\alpha = 0.05$, (c) $\alpha = 0.10$, and (d) $\alpha = 0.25$.

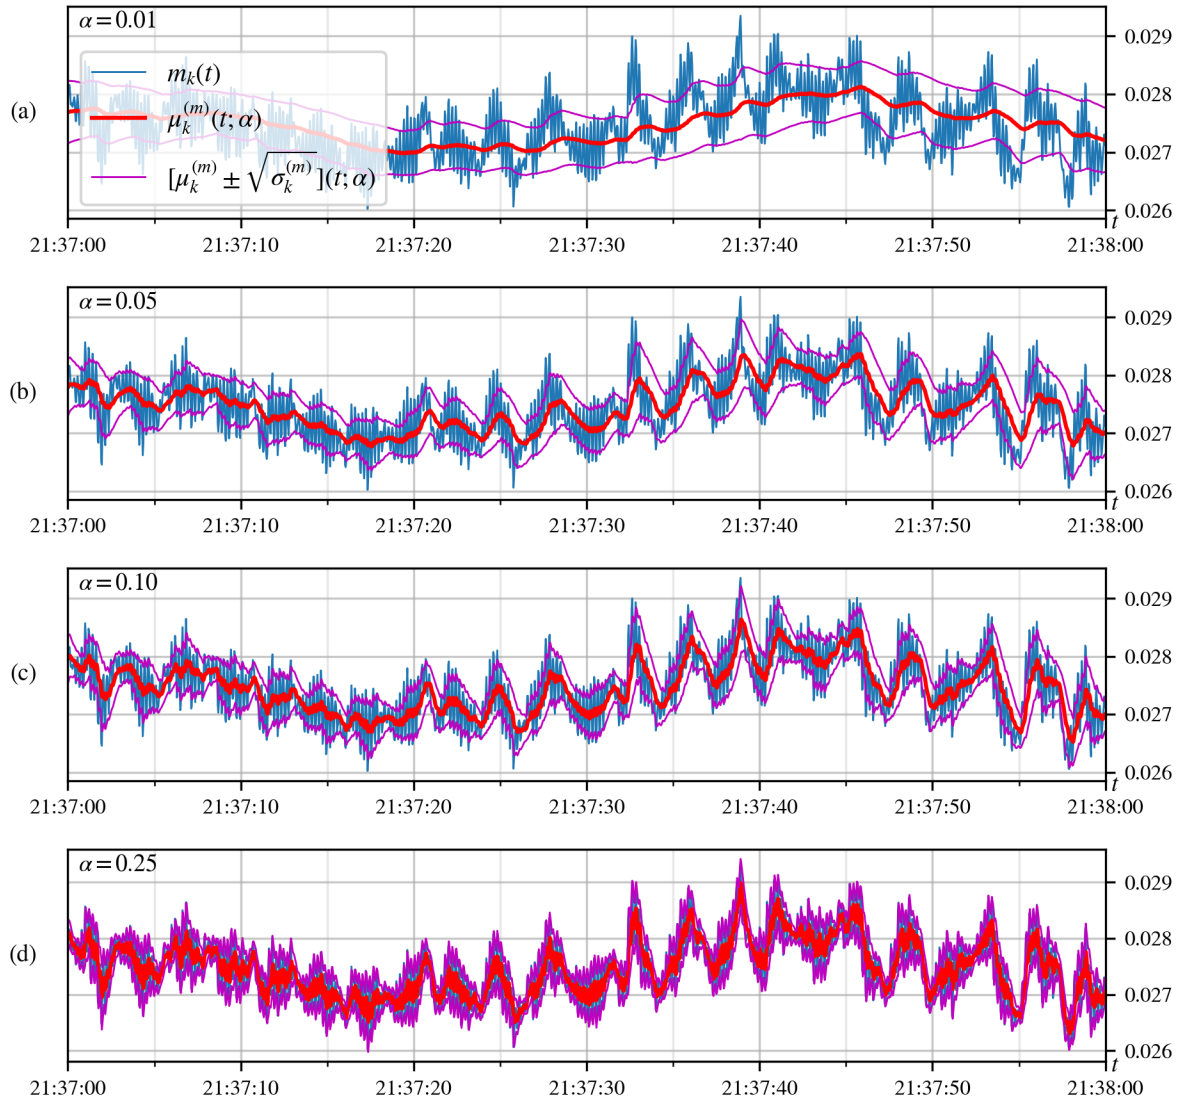

Figure 4.3: Voltage phase angle measurements with respect to a reference bus (in blue) of PMU $k = 156$ during one minute, between 9:37 and 9:38 PM on July 30, 2013. Moving average (in red) and a band of one standard deviation (in purple) are also shown for (a) $\alpha = 0.01$, (b) $\alpha = 0.05$, (c) $\alpha = 0.10$, and (d) $\alpha = 0.25$.

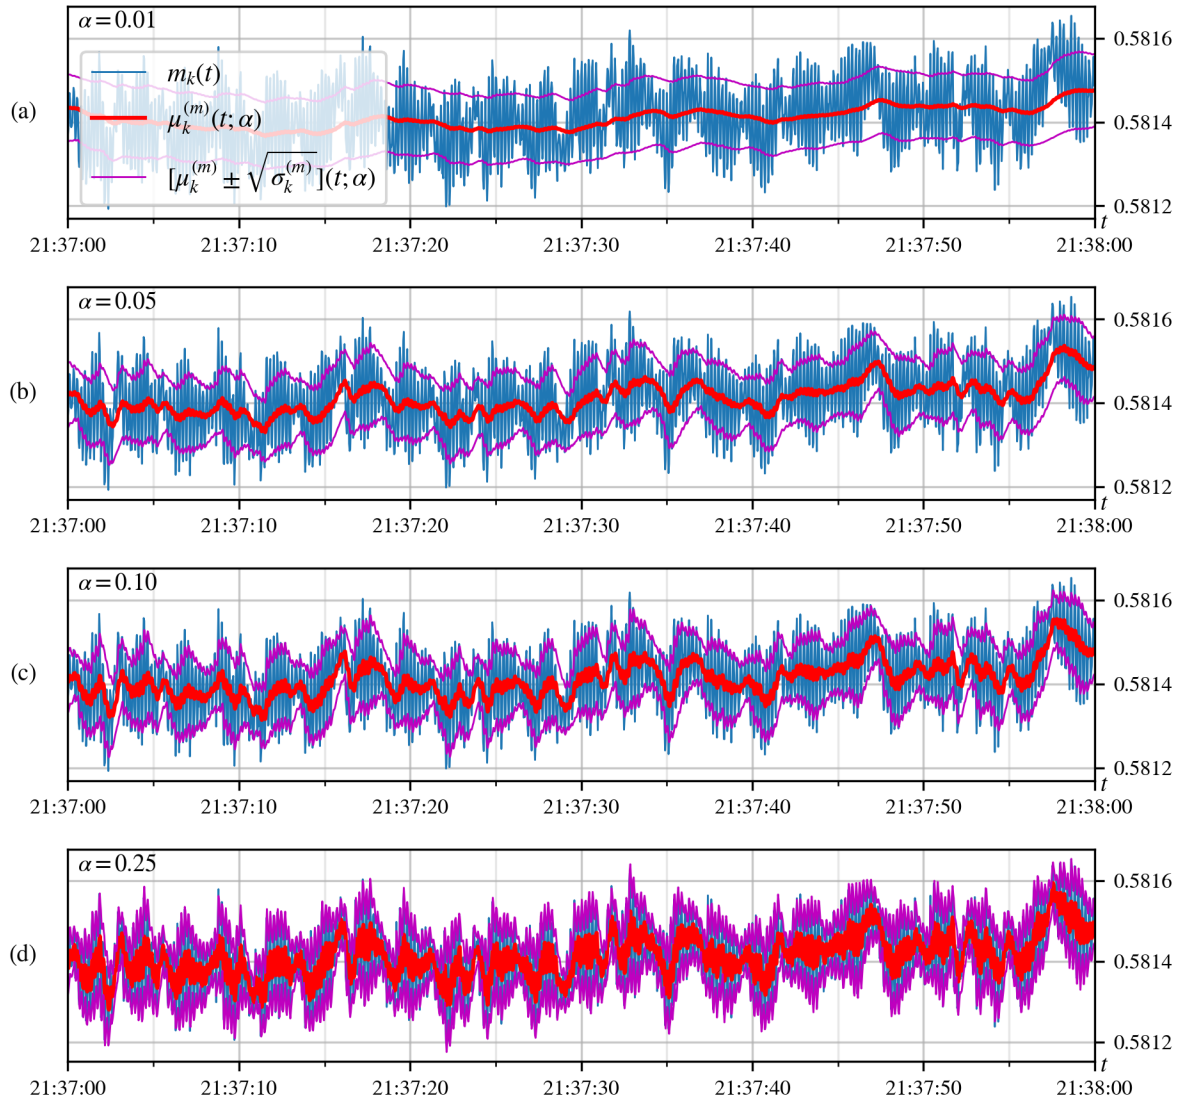

Figure 4.4: Voltage magnitude measurements with respect to bus nominal value (in blue) of PMU $k = 156$ during one minute, between 9:37 and 9:38 PM on July 30, 2013. Moving average (in red) and a band of one standard deviation (in purple) are also shown for (a) $\alpha = 0.01$, (b) $\alpha = 0.05$, (c) $\alpha = 0.10$, and (d) $\alpha = 0.25$.

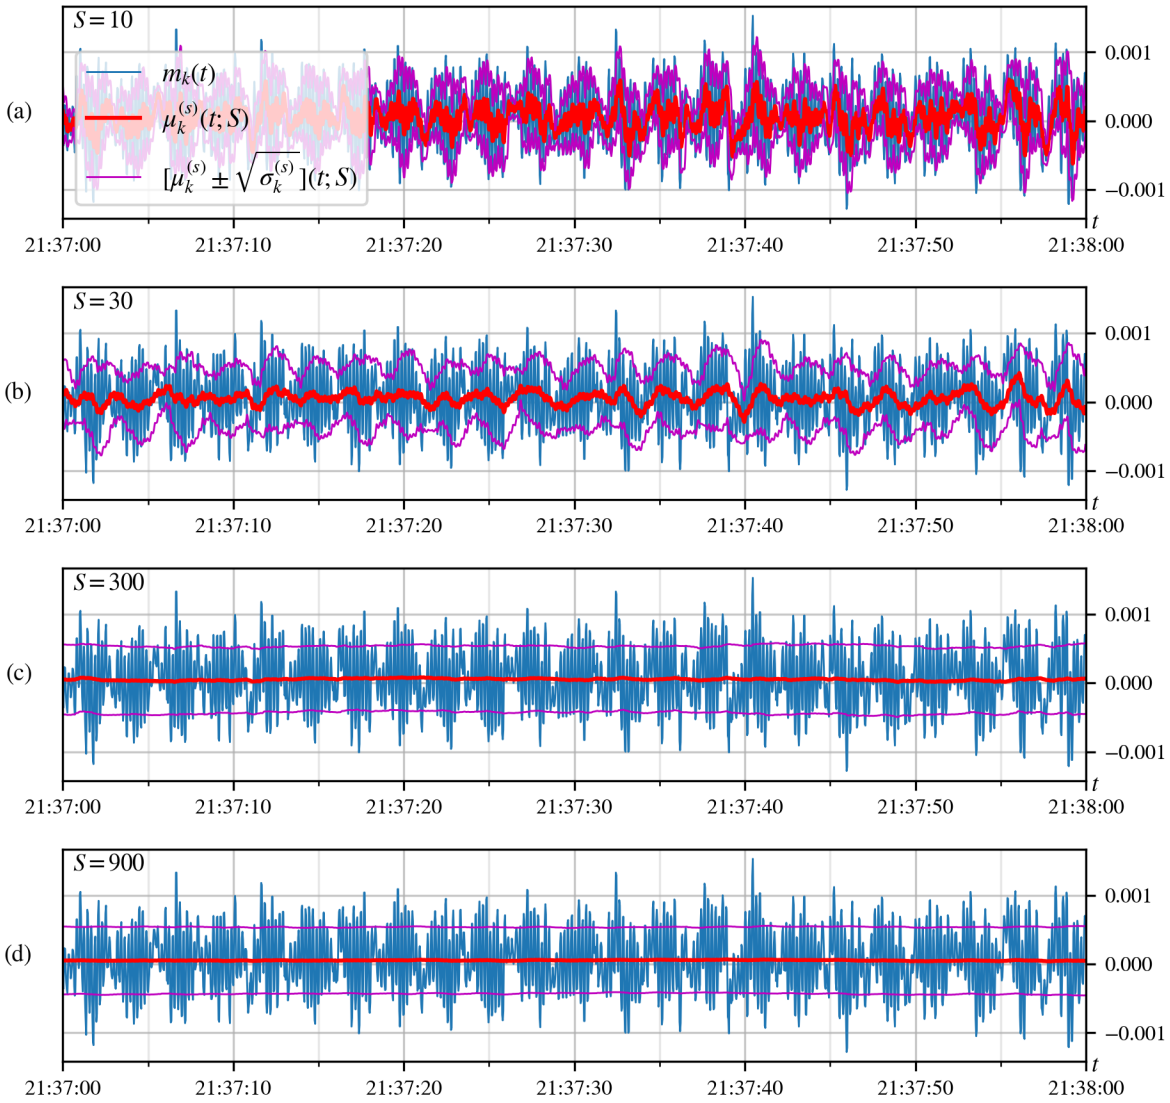

Figure 4.5: Frequency measurements with respect to a reference bus (in blue) of PMU $k = 156$ during one minute, between 9:37 and 9:38 PM on July 30, 2013. Average over sliding time horizon (in red) and a band of one standard deviation (in purple) are also shown for (a) $S = 10$, (b) $S = 30$, (c) $S = 300$, and (d) $S = 900$; corresponding to the number of samples made over 1/3, 1, 10, and 30 seconds, respectively.

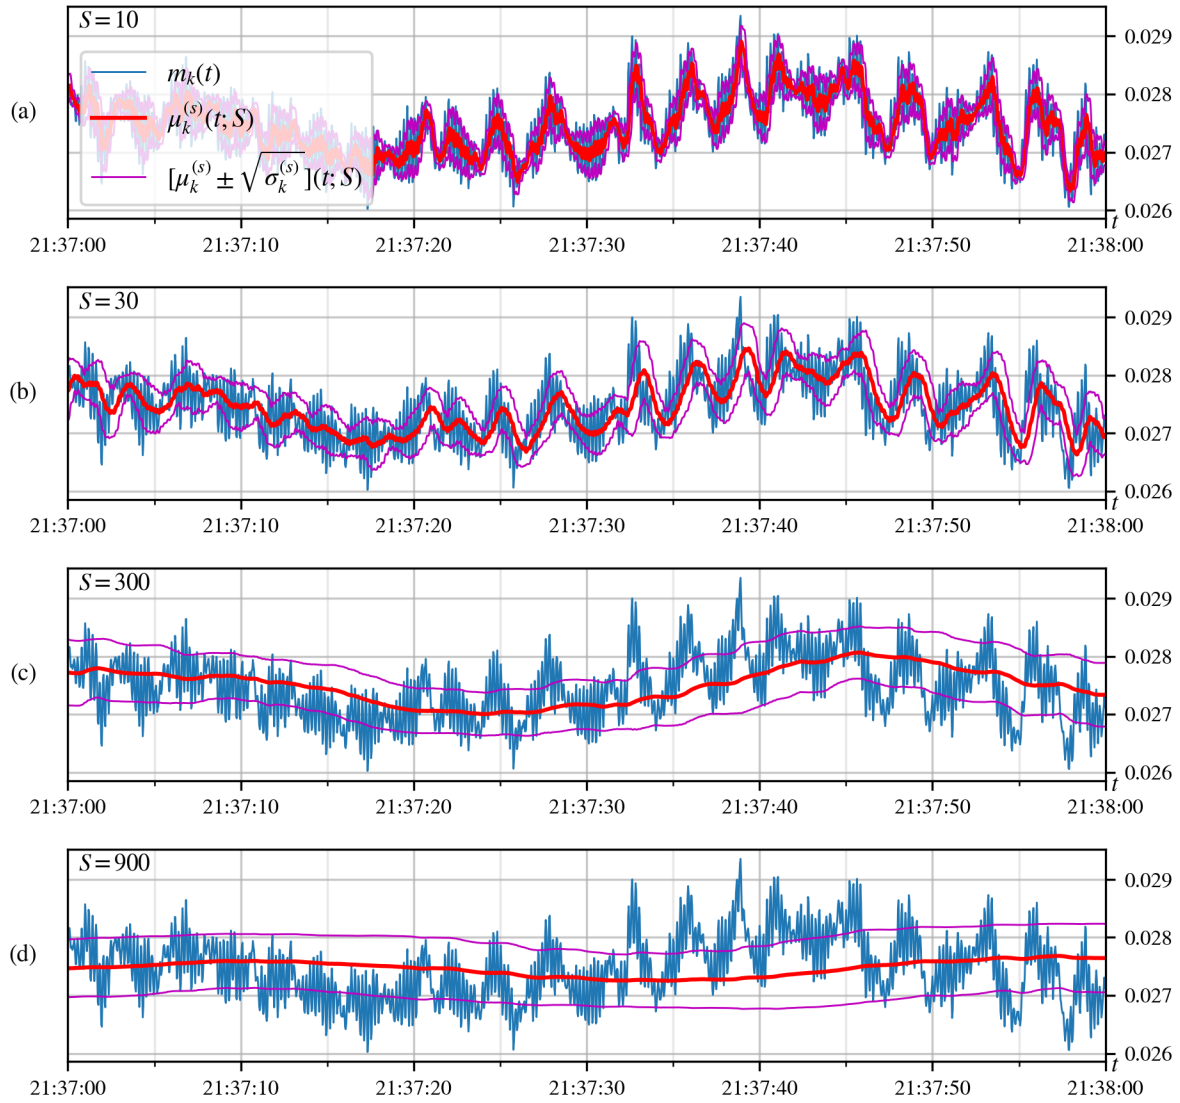

Figure 4.6: Voltage phase angle measurements with respect to a reference bus (in blue) of PMU $k = 156$ during one minute, between 9:37 and 9:38 PM on July 30, 2013. Average over sliding time horizon (in red) and a band of one standard deviation (in purple) are also shown for (a) $S = 10$, (b) $S = 30$, (c) $S = 300$, and (d) $S = 900$; corresponding to the number of samples made over 1/3, 1, 10, and 30 seconds, respectively.

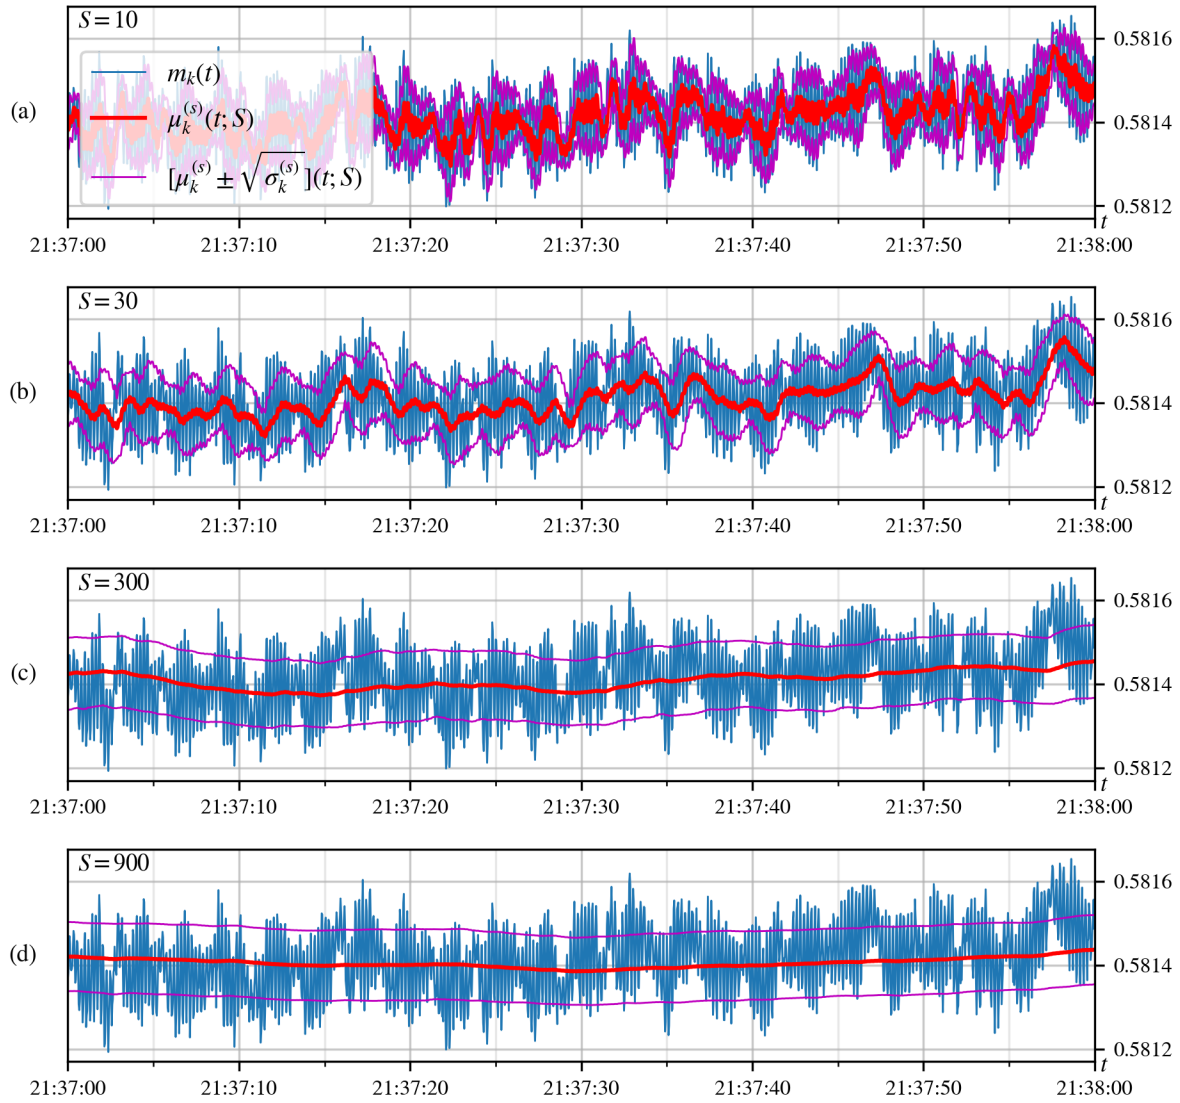

Figure 4.7: Voltage magnitude measurements with respect to bus nominal value (in blue) of PMU $k = 156$ during one minute, between 9:37 and 9:38 PM on July 30, 2013. Average over sliding time horizon (in red) and a band of one standard deviation (in purple) are also shown for (a) $S = 10$, (b) $S = 30$, (c) $S = 300$, and (d) $S = 900$; corresponding to the number of samples made over 1/3, 1, 10, and 30 seconds, respectively.

4.4 Fourier Filtering

In the following analysis we assume that Q is an even integer. Let $\mathcal{F}[m_k(\cdot)] \in \mathbb{C}^Q$ be the discrete Fourier transform of $m_k(\cdot)$, where $\mathcal{F}[m_k(\cdot)]_q$ is the amplitude corresponding to frequency ω_q, $q \in \{0, \dots, Q-1\}$. Since the frequency of the readings is 30 Hz, we can represent the frequency domain as Q equidistant points between 0 and 30 Hz. In other words,

$$\forall q \in \{0, \dots, Q-1\}, \qquad \mathcal{F}[m_k(\cdot)]_q \doteq \sum_{t=0}^{Q-1} m_k(t) \exp\left\{-2\pi j \frac{q\omega_t}{30}\right\},$$
$$\text{where} \quad \omega_t \doteq 30 \cdot \frac{t}{Q}.$$

Given a vector $x = (x_q)_{q \in \{0,\dots,Q-1\}}$, the inverse Fourier transform is defined as

$$\forall t \in \{0, \dots, Q-1\}, \qquad \mathcal{F}^{-1}[x]_t \doteq \frac{1}{Q} \sum_{q=0}^{Q-1} x_q \exp\left\{2\pi j \frac{t\omega_q}{30}\right\}$$

As it is well known, we have the following

Observation 16 ([71]). *The inverse Fourier transform of the Fourier transform recovers the original time series, that is*

$$\forall t \in \{0, \dots, Q-1\}, \qquad m_k(t) = \mathcal{F}^{-1}\left[\mathcal{F}[m_k(\cdot)]\right].$$

These relations give a one-to-one correspondence between the original time series (characterized in the time domain t) and its Fourier transform series (characterized in the frequency domain ω_q). Moreover, given the periodicity of the sinusoidal functions, we can represent the time series as

$$m_k(t) = \frac{1}{Q} \sum_{q=-Q/2}^{Q/2-1} \mathcal{F}^{\text{shift}}[m_k(\cdot)]_q \exp\left\{2\pi j \frac{t\omega_q}{30}\right\}, \tag{4.9}$$

where the frequencies ω_q spans from -15 and 15 Hz, and

$$\mathcal{F}^{\text{shift}}[\cdot]_q = \begin{cases} \mathcal{F}[\cdot]_{Q+q}, & \text{if } q < 0, \\ \mathcal{F}[\cdot]_q, & \text{if } q \geq 0. \end{cases}$$

$\mathcal{F}^{\text{shift}}$ represents the shifted Fourier transform, placing the zero-frequency component to the center of the spectrum. This representation of the Fourier transform is widely used in signal processing since the symmetry with respect to zero ease its visualization and the application of different filters.

Observation 17. *$(x(t))_{t\in\{0,\dots,Q-1\}}$ is a real-valued time series, if and only if, $\mathcal{F}[x]_0$ and $\mathcal{F}[x]_{Q/2}$ are real and for any $q \in \{1,\dots,\frac{Q}{2}-1\}$, $\mathcal{F}[x]_q = \overline{\mathcal{F}[x]_{Q-q}}$.*

Proof. ($\Rightarrow$) By definition of the Fourier transform, it is easy to see that $\mathcal{F}[x]_0$ and $\mathcal{F}[x]_{Q/2}$ are real when x is real. Let $q > 0$,

$$\begin{aligned}\mathcal{F}[x]_q &= \frac{1}{Q}\sum_{t=0}^{Q-1} m_k(t)\exp\left\{-2\pi j\frac{q\omega_t}{30}\right\}\\ &= \frac{1}{Q}\sum_{t=0}^{Q-1} m_k(t)\exp\left\{+2\pi j\frac{(Q-q)\omega_t}{30}\right\} = \overline{\mathcal{F}[x]_{Q-q}},\end{aligned}$$

The last equality holds, since by definition $\frac{Q\omega_t}{30} = t$, and therefore $\exp\{2\pi j\frac{Q\omega_t}{30}\} = 1$.
($\Leftarrow$) Let $t \in \{0,\dots,Q-1\}$, using Observation 16 and the fact that $\exp\{2\pi j\frac{t\omega_{Q-q}}{30}\} = \exp\{2\pi j\frac{t(Q-q)}{Q}\} = \exp\{-2\pi j\frac{tq}{Q}\} = \exp\{-2\pi j\frac{t\omega_q}{30}\}$,

$$\begin{aligned}Q\cdot x(t) &= \mathcal{F}[x]_0 + \mathcal{F}[x]_{Q/2}\exp\left\{2\pi j\frac{t\omega_{Q/2}}{30}\right\}\\ &\quad + \sum_{q=1}^{Q/2-1}\left[\mathcal{F}[x]_q\exp\left\{2\pi j\frac{t\omega_q}{30}\right\} + \mathcal{F}[x]_{Q-q}\exp\left\{2\pi j\frac{t\omega_{Q-q}}{30}\right\}\right]\\ &= \mathcal{F}[x]_0 + \mathcal{F}[x]_{Q/2}\exp\{\pi tj\}\\ &\quad + \sum_{q=1}^{Q/2-1}\left[\overline{\mathcal{F}[x]_{Q-q}}\exp\left\{2\pi j\frac{t\omega_q}{30}\right\} + \mathcal{F}[x]_{Q-q}\exp\left\{-2\pi j\frac{t\omega_q}{30}\right\}\right].\end{aligned}$$

Since the two first terms in the last expression are real, and the summation goes over terms that add a quantity and its conjugate, we conclude that $Q\cdot x(t)$ is real. □

The property that a time series satisfies $x_q = \overline{x_{Q-q}}$ will be called *conjugate-symmetry*.

Corollary 18. *If $(x(t))_{t\in\{0,\dots,Q-1\}}$ is a real-valued time series, then $\left|\mathcal{F}^{\text{shift}}[x]_q\right| = |\mathcal{F}[x]_q| = |\mathcal{F}[x]_{Q-q}| = \left|\mathcal{F}^{\text{shift}}[x]_{-q}\right|$ for any $q \in \{1,\dots,\frac{Q}{2}-1\}$.*

Remark. An analogous result to Observation 17 can be obtain for the inverse Fourier transform. Therefore, we conclude that a symmetric real-valued —and, therefore, conjugate-symmetric— time series has a Fourier transform that is also real-valued and (conjugate-) symmetric.

Figure 4.8 shows the absolute value of the Fourier transform for a particular PMU during a 15-minute quiet period. We observe predominant peaks near 5 Hz in the frequency domain of the signal, this particularity repeats over different PMUs and different days and hours. Figure 4.9 shows the absolute value of the Fourier transform of the complex voltage of the same time series as Figure 4.8.

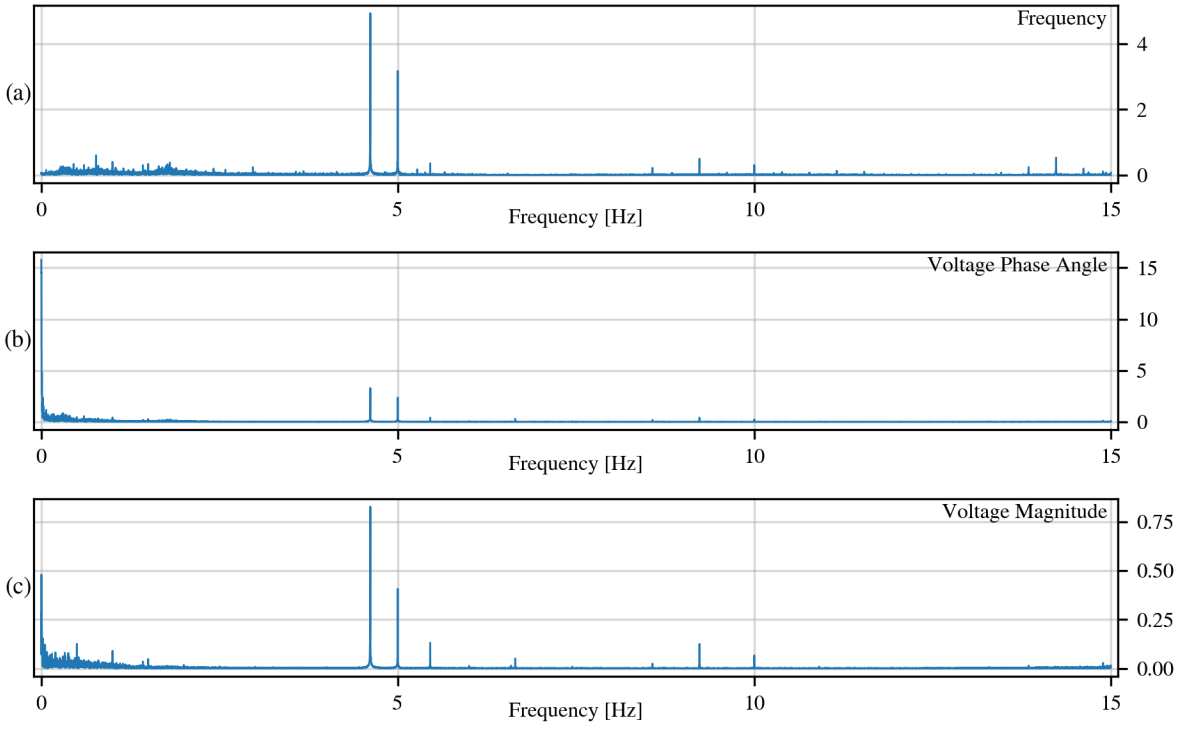

Figure 4.8: Absolute value of the Fourier transform over the positive frequency domain (0 to 15 Hz) of (a) frequency, (b) voltage phase angle, and (c) voltage magnitude of PMU $k = 156$ over quiet period #10 (see Table 4.1). Mean have been subtracted from the raw signal —to prevent a high peaks at frequency 0.

Consider a vector $\tilde{\varphi} = \left\{\tilde{\varphi}_q : q \in \{-\frac{Q}{2}, \ldots, \frac{Q}{2} - 1\}\right\} \in \mathbb{C}^Q$.

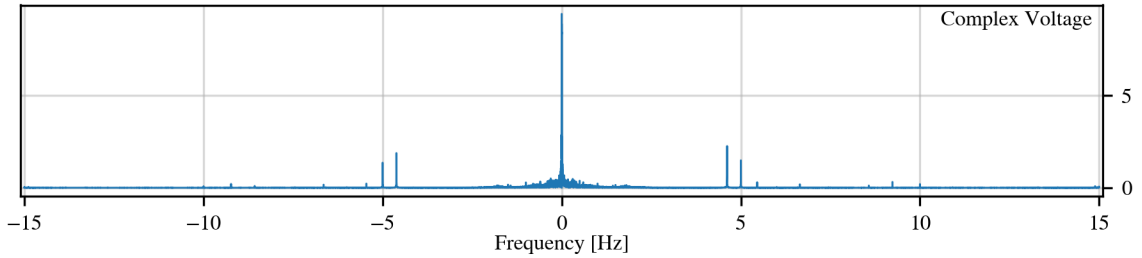

Figure 4.9: Absolute value of the Fourier transform of the complex voltage of PMU $k = 156$ over quiet period #10. Again, mean have been subtracted from the raw signal.

Definition 19. *We denote by* $(f_k[\tilde{\varphi}](t))_{t\in\{0,\dots,Q-1\}}$ *the* $\tilde{\varphi}$*-filtered time series of* $m_k(\cdot)$ *that is defined as the vector whose Fourier transform is the result of multiplying component-wise each of the terms of* $\tilde{\varphi}$ *and the shifted Fourier transform of* $m_k(\cdot)$*.*

In mathematical terms, the previous definition can be written as

$$\forall q \in \{0,\dots,Q-1\}, \qquad \mathcal{F}[f_k[\tilde{\varphi}](\cdot)]_q \;=\; \tilde{\varphi}_q \cdot \mathcal{F}^{\text{shift}}[m_k(\cdot)]_q.$$

We will use the operator $\odot$ to denote the component-wise product of two vectors, therefore we can write $\mathcal{F}[f_k[\tilde{\varphi}](\cdot)] \;=\; \tilde{\varphi} \odot \mathcal{F}^{\text{shift}}[m_k(\cdot)]$. Equivalently, we have:

$$\forall t \in \{0,\dots,Q-1\}, \qquad f_k[\tilde{\varphi}](t) \;\doteq\; \frac{1}{Q}\sum_{q=-Q/2}^{Q/2-1} \tilde{\varphi}_q \cdot \mathcal{F}^{\text{shift}}[m_k(\cdot)]_q \exp\left\{2\pi j \frac{t\omega_q}{30}\right\}. \tag{4.10}$$

If we extend the definition of the vector $\tilde{\varphi}$, in such a way that it repeats its values periodically, meaning that for any integer n, $\tilde{\varphi}_{nQ+q} = \tilde{\varphi}_q$ for all $q \in \{-\frac{Q}{2},\dots,\frac{Q}{2}-1\}$; we can equivalently write (4.10) as

$$f_k[\tilde{\varphi}](t) \;=\; \frac{1}{Q}\sum_{q=0}^{Q-1} \tilde{\varphi}_q \cdot \mathcal{F}[m_k(\cdot)]_q \exp\left\{2\pi j \frac{t\omega_q}{30}\right\}. \tag{4.11}$$

Let us denote $\varphi \doteq \mathcal{F}^{-1}[\tilde{\varphi}]$ the inverse Fourier transform of $\tilde{\varphi}$. Then,

Lemma 20 ([71]). $\forall t \in \{0,\dots,Q-1\}$, $f_k[\tilde{\varphi}](t) \;=\; \sum_{s=0}^{Q-1} m_k(s) \cdot \varphi_{(t-s)}$.

Therefore, we can express the filtered series as in Definition 19 or as a convolution between the original series $m_k(\cdot)$ and φ.

Analogously to equations (4.5)–(4.8), we can compute (averaging over a sliding time horizon) the zero-mean data vector and the re-scaled zero-mean data vector for $f_k[\tilde{\varphi}](\cdot)$ instead of $m_k(\cdot)$, we will call them $\bar{f}_k^{(s)}[\tilde{\varphi}](\cdot; S)$ and $\hat{f}_k^{(s)}[\tilde{\varphi}](\cdot; S)$, respectively, when the average is made over the last S measurements.

We will define next several filter vectors and compare the time series that result after applying the corresponding Fourier filters.

4.4.1 Cutting off high frequency components

We want to understand and visualize how the filtered data looks like when high frequencies are suppressed. For that purpose, we will consider the following family of cut-off (CO) filter vectors:

$$\forall q \in \{-\tfrac{Q}{2}, \dots, \tfrac{Q}{2} - 1\}, \qquad \tilde{\varphi}_q^{\mathrm{CO}\lambda} \;=\; (\chi_{[-\lambda,\lambda]})_q \;\doteq\; \begin{cases} 1, & \text{if } |\omega_q| \leq \lambda, \\ 0, & \text{otherwise,} \end{cases} \tag{4.12}$$

where $\lambda \in (0, 15)$ is the cut-off parameter.

Figure 4.10 shows the filter vector $\tilde{\varphi}^{\mathrm{CO}\lambda}$ and its inverse Fourier transform $\varphi^{\mathrm{CO}\lambda}$, for $\lambda = 4$ Hz. Figure 4.11 shows (on plots on the left) the original time series for frequency, voltage phase angle and voltage magnitude of PMU $k = 156$ across 10 seconds of measurements (300 samples). The plots on the right show the resulting filtered time series when the filter vector $\tilde{\varphi}^{\mathrm{CO}4}$ is applied. As a remainder, the filter has been applied to the frequency domain of the Fourier filter obtained on an interval of 15 minutes.

We also show that when we cut off frequencies higher than the peaks around 5 Hz that we see in Figures 4.8 and 4.9, for example, cutting off frequencies larger that 6 Hz, we obtain almost no change between the original and the filtered time series. Figure 4.12 show the cut-off filter vector $\tilde{\varphi}^{\mathrm{CO}6}$ and the filtering results are shown in Figure 4.13.

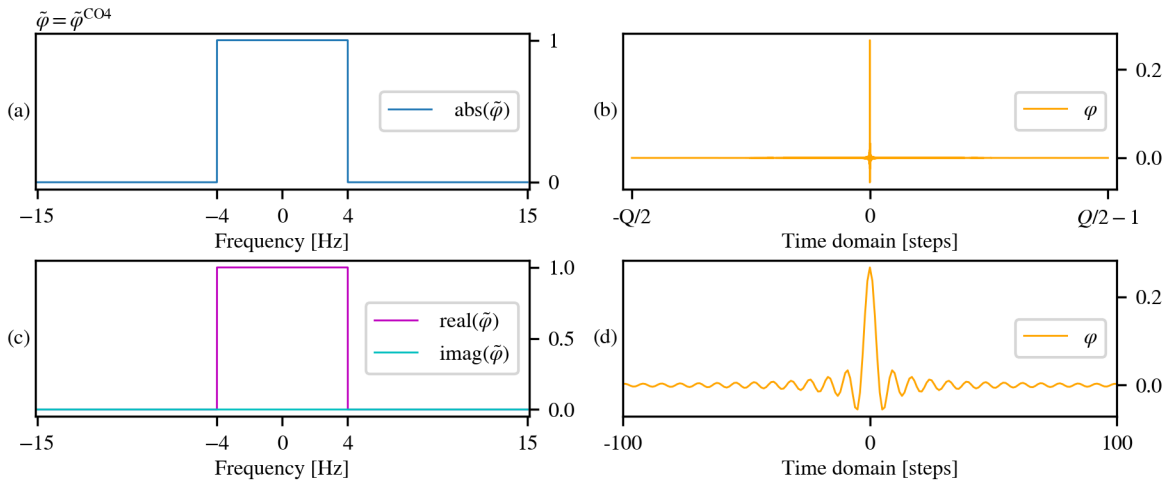

Figure 4.10: (a) Absolute value and (c) real and imaginary parts of filter vector $\tilde{\varphi}^{CO4}$; (b) and (d) show the component-wise absolute value of the inverse Fourier transform of $\tilde{\varphi}^{CO4}$.

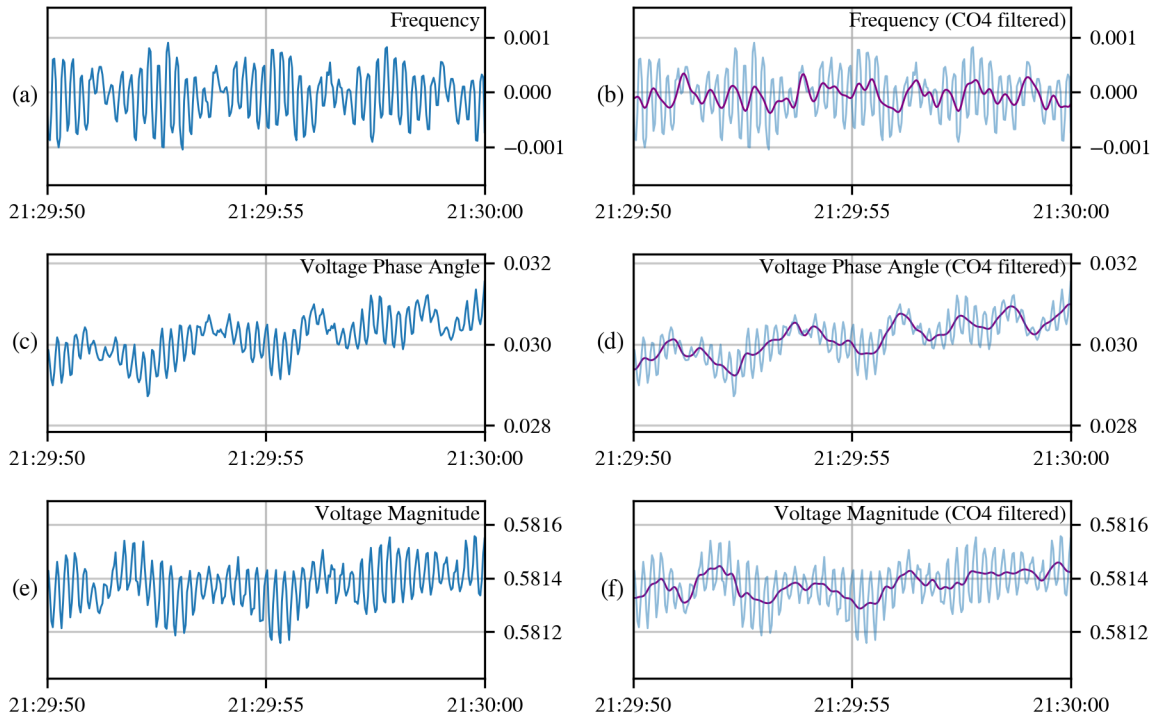

Figure 4.11: Original sampling (in blue) for (a) frequency, (c) voltage phase angle, and (e) voltage magnitude across 10 seconds on period #10 (see Table 4.1). Corresponding filtered series are shown (in purple) in plots (b), (d), and (f) using filtering vector $\tilde{\varphi}^{CO4}$.

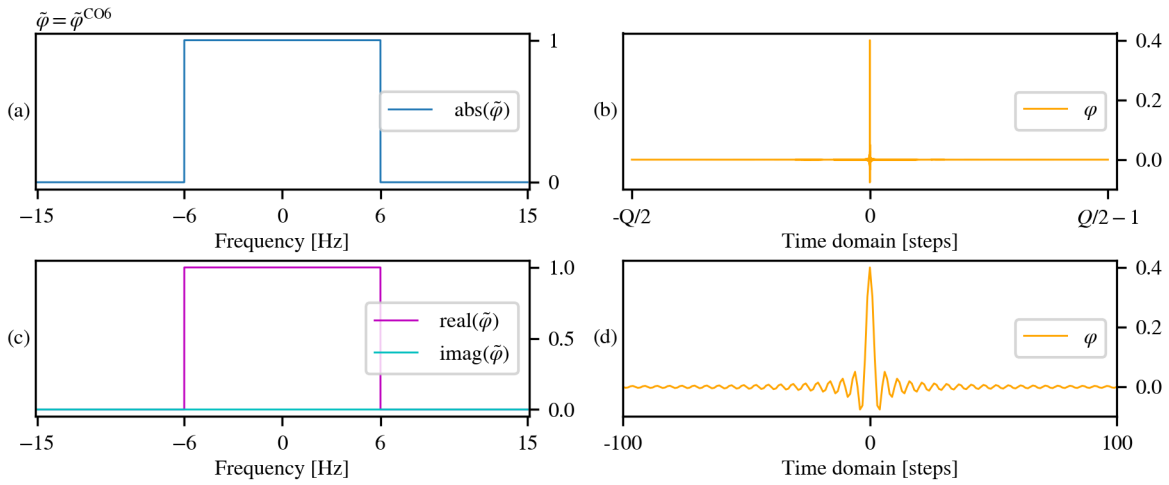

Figure 4.12: (a) Absolute value and (c) real and imaginary parts of filter vector $\tilde{\varphi}^{CO6}$; (b) and (d) show the component-wise absolute value of the inverse Fourier transform of $\tilde{\varphi}^{CO6}$.

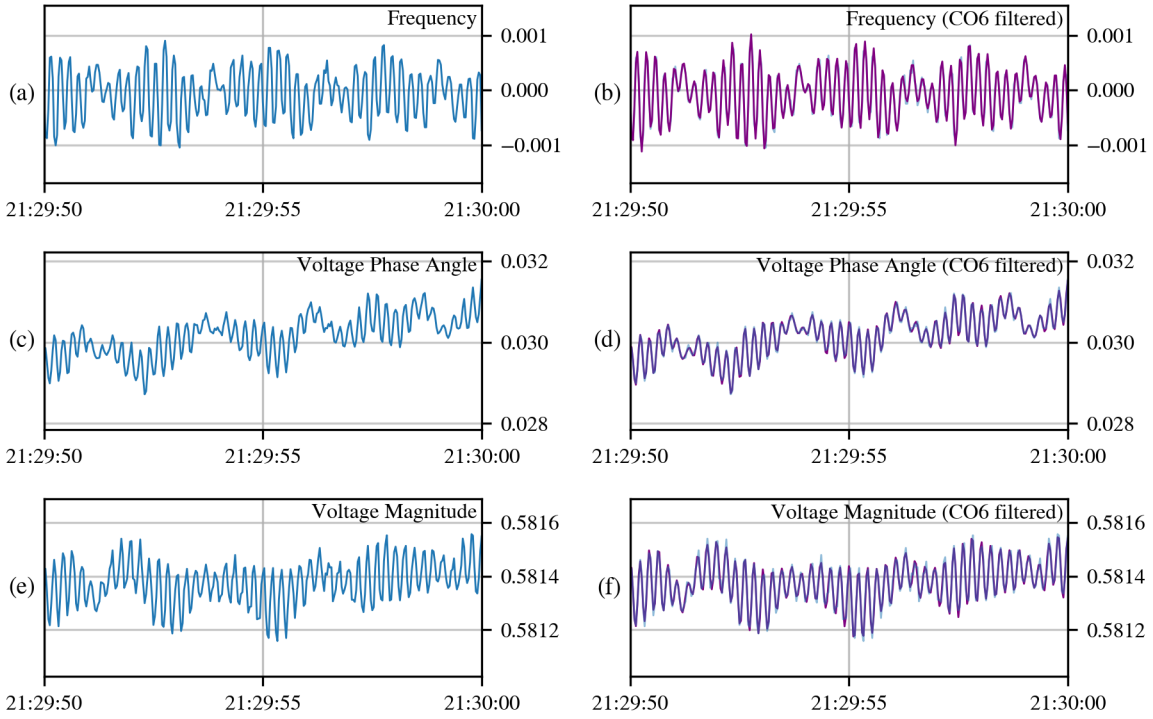

Figure 4.13: Original sampling (in blue) for (a) frequency, (c) voltage phase angle, and (e) voltage magnitude across 10 seconds on period #10 (see Table 4.1). Corresponding filtered series are shown (in purple) in plots (b), (d), and (f) using filtering vector $\tilde{\varphi}^{CO6}$.

More sophisticated filters are used in signal processing, so that the cut-off process can be made in a online and smooth fashion. The cut-off vector presented above is an example of a low-pass filter (LPF), meaning that the filter passes lower frequencies than a selected cut-off frequency.

In the following examples we present LPF vectors that are first described by its inverse Fourier transform $\varphi^{\mathrm{LPF}\lambda}$ —i.e. the coefficients that will participate in the convolution to get the filtered series— for cut-off frequency λ. In this filtering process, the convolution is made with a small number (compared with Q) of measurements, the rest of coefficients are set to be zero; this property allows to have an online computation of the filtered time series. These filters are called *finite impulse response* (FIR) systems [71].

On the other hand, the *infinite impulse response* (IIR) systems consist of infinite vectors of non-zero coefficients that are convoluted with an infinite time series (in our case, we might think about extending our original time series by repeating it periodically out of the $[0, Q-1]$ interval). For illustrative purposes, consider the cut-off filter vectors (4.12) in the continuum set-up:

$$\forall \omega \in [-15, 15], \qquad H_d(\omega) \;=\; \begin{cases} 1, & \text{if } |\omega| \leq \lambda, \\ 0, & \text{otherwise,} \end{cases} \tag{4.13}$$

this is the ideal *desired* frequency response and can be represented as

$$H_d(\omega) \;=\; \sum_{t=-\infty}^{\infty} h_d(t) \exp\left\{-2\pi j \frac{\omega t}{30}\right\}, \tag{4.14}$$

where h_d is the corresponding IIR sequence, that expressed in terms of H_d becomes

$$\forall t \in \mathbb{Z}, \qquad h_d(t) \;=\; \frac{1}{30} \int_{-15}^{15} H_d(\omega) \exp\left\{2\pi j \frac{t\omega}{30}\right\} d\omega. \tag{4.15}$$

Figure 4.14 shows the desired frequency response H_d for $\lambda = 5$ Hz and its partial IIR sequence. As expected, the IIR vector assimilates the ones computed above for $\varphi^{\mathrm{CO}\lambda}$ (see Figures 4.12d and 4.10d).

Given the impractibality of convoluting a times series by an infinite vector, the computation is usually truncated or, equivalently, the IIR sequence is truncated to a chosen bounded window and setting equal to zero the terms that fall out of the selected window.

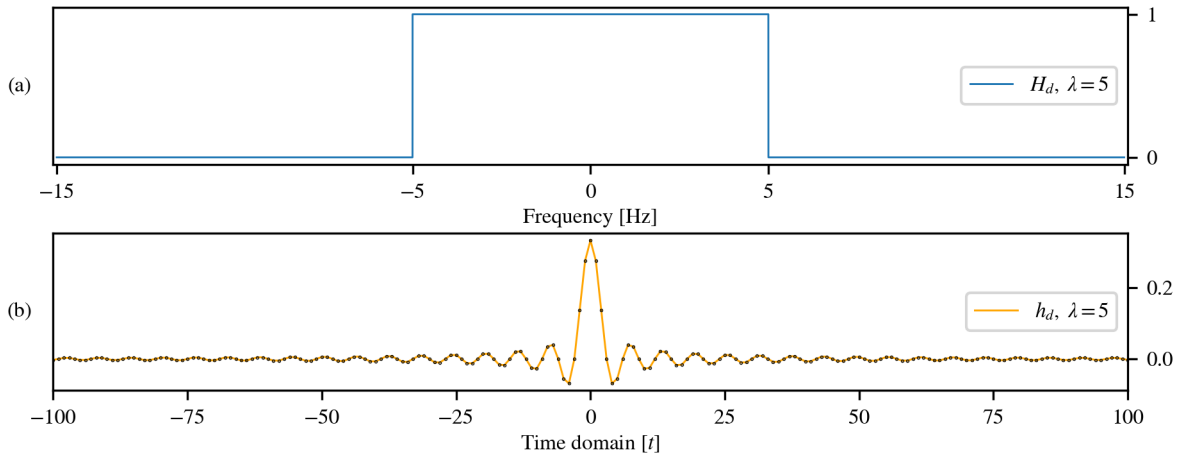

Figure 4.14: (a) Desired frequency response H_d for cut-off frequency $\lambda = 5$ Hz and (b) its infinite impulse response (showing for t between -100 and 100).

As explained in [71], one technique of designing FIR filters is by *windowing*. That is, given a finite-duration window sequence w, where

$$w(t) \;=\; 0 \quad \text{for } t < -M_1 \text{ or } t > M_2, \qquad\qquad (\text{with } M_1, M_2 \in \mathbb{N}),$$

a FIR filter is obtained by the componen-wise product between h_d and w:

$$\forall t \in \mathbb{Z}, \qquad h(t) \;\doteq\; h_d(t)w(t), \tag{4.16}$$

(or, equivalently, $h = h_d \odot w$). Since filtering the original time series $m_k(\cdot)$ by the filter vector $h(\cdot)$ corresponds to the convolution in the time domain of this two sequences, in order to obtain the filtered element at some time t we need M_1 future measurements. Therefore, we can obtain an online filtered series with a delay of M_1 time steps.

Different type of windows have been proposed in the literature and give appropriate results depending on the nature and the pyshics behind of the time series. The simplest one if the rectangular window

$$w_1(t) \;\doteq\; \begin{cases} 1, & \text{if } -M_1 \leq t \leq M_2, \\ 0, & \text{otherwise,} \end{cases} \tag{4.17}$$

that simply truncates the IIR to the $[-M_1, M_2]$ window. The average over a sliding time horizon defined in (4.5) can be computed as a filtered time series with FIR filter $h = h_d \odot w$, where $h_d(t) = 1$ for all t and w is a rectangular window with $M_1 = 0$ and $M_2 = S - 1$. In the following examples, we will use the Hann window [44] defined as

$$w_2(t) \doteq \begin{cases} \frac{1}{2}\left[1 - \cos\left(\frac{2\pi(t+M)}{2M}\right)\right] = \sin^2\left(\frac{\pi(t+M)}{2M}\right), & \text{if } -M \leq t \leq M, \\ 0, & \text{otherwise,} \end{cases} \tag{4.18}$$

for $M \in \mathbb{N}$, since it has given us satisfactory results. Figure 4.15 shows the positive coefficients of a rectangular window and a Hann window.

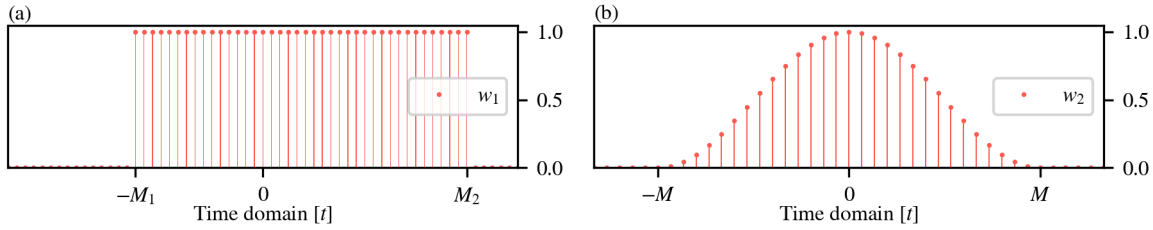

Figure 4.15: (a) Rectangular window and (b) Hann window.

By increasing M we obtaing a filter that is closer to the desired frequency response, but at the same time, we increase the delay of the online filtering process. Therefore, we must look for a small parameter M that gives useful filtered time series.

Figure 4.16 shows the low-pass filter for $\lambda = 4$ Hz described by 101 coefficients (by setting $M = 50$ in (4.18)), shown as black dots in plot (d). The Fourier transform of the FIR filter $h = \varphi^{\text{LPF4}}$ results as a smoother version of the one obtained by φ^{CO4} (compare Figures 4.10a and 4.16a).

Finally, we show that the amount of coefficients used to describe the LPF vector is related with the slope of the vector from the region of frequencies that are kept and the ones that are cut; the more coefficients the steepest the slope is. Figures 4.18 to 4.29 show the LPF vector $\varphi^{\text{LPF}\lambda}$ with 31 coefficients (by setting $M = 15$ in (4.18)) and the resulting filtered time series for $\lambda = 1$, 2, 3, 4, 5, and 6 Hz.

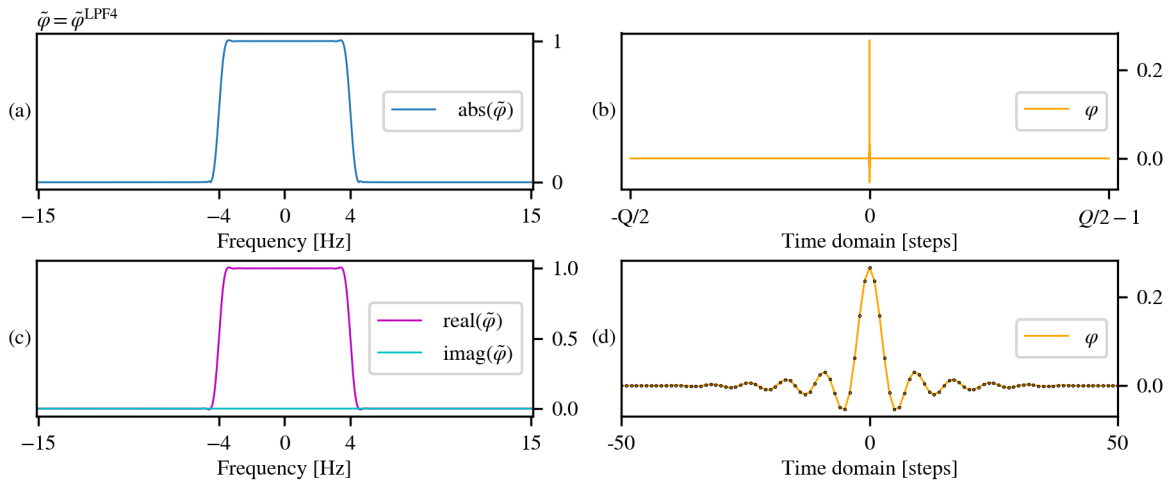

Figure 4.16: (a) Absolute value, (c) real and imaginary parts of filter vector $\tilde{\varphi}^{\mathrm{LPF4}}$ ($M = 50$); (b) and (d) show the component-wise absolute value of the inverse Fourier transform of $\tilde{\varphi}^{\mathrm{LPF4}}$.

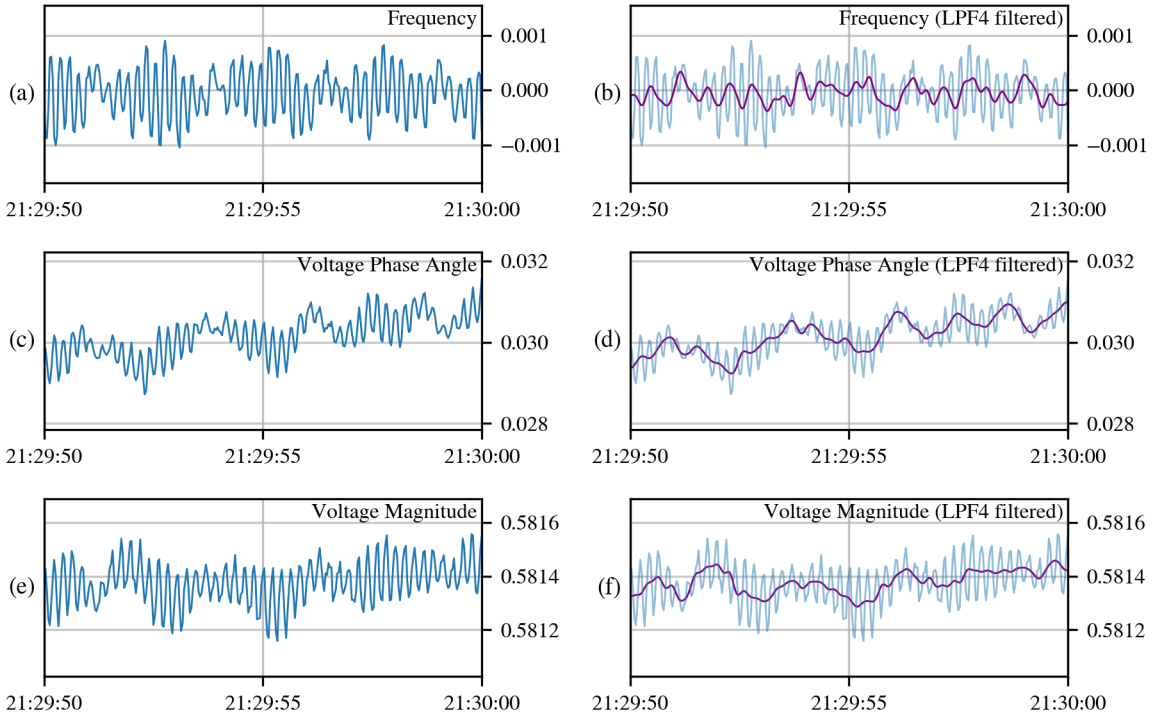

Figure 4.17: Original sampling (in blue) for (a) frequency, (c) voltage phase angle, and (e) voltage magnitude across 10 seconds on period #10 (see Table 4.1). Corresponding filtered series are shown (in purple) in plots (b), (d), and (f) using filtering vector $\tilde{\varphi}^{\mathrm{LPF4}}$.

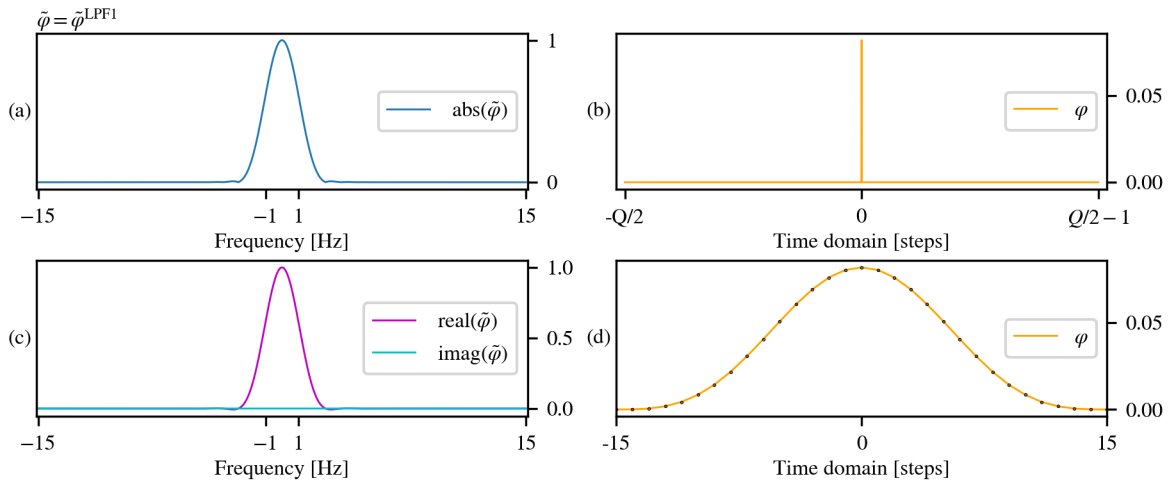

Figure 4.18: (a) Absolute value and (c) real and imaginary parts of filter vector $\tilde{\varphi}^{\mathrm{LPF1}}$ ($M = 15$); (b) and (d) show the component-wise absolute value of the inverse Fourier transform of $\tilde{\varphi}^{\mathrm{LPF1}}$.

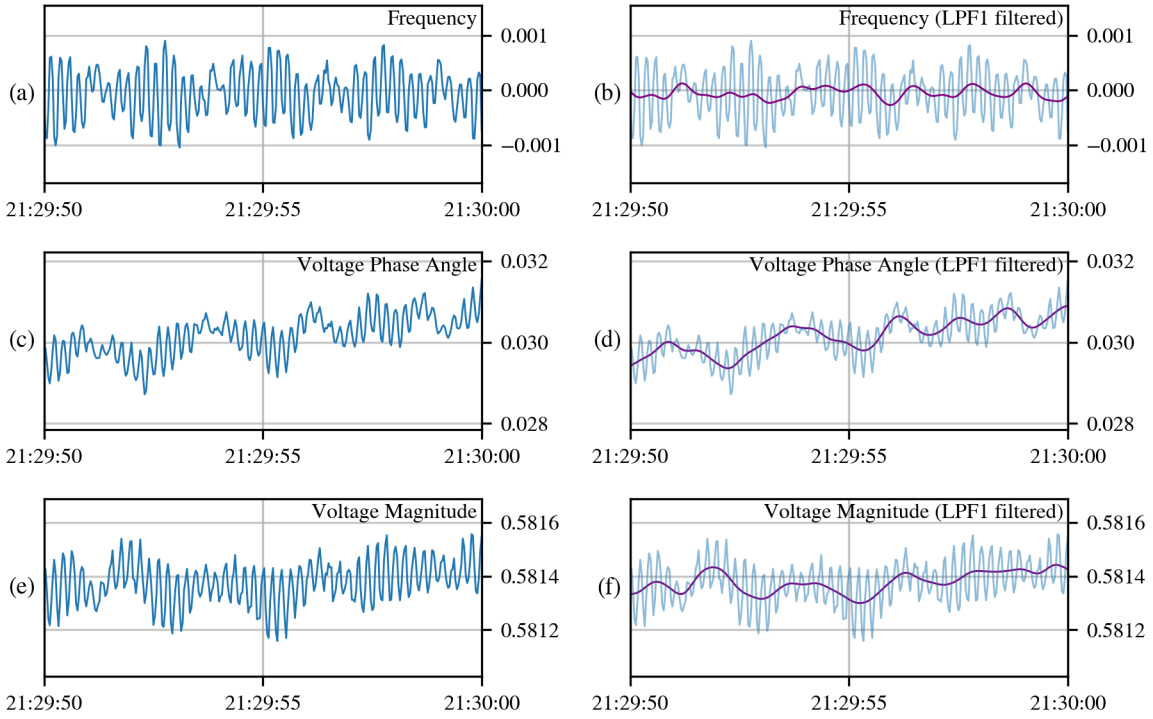

Figure 4.19: Original sampling (in blue) for (a) frequency, (c) voltage phase angle, and (e) voltage magnitude across 10 seconds on period #10 (see Table 4.1). Corresponding filtered series are shown (in purple) in plots (b), (d), and (f) using filtering vector $\tilde{\varphi}^{\mathrm{LPF1}}$.

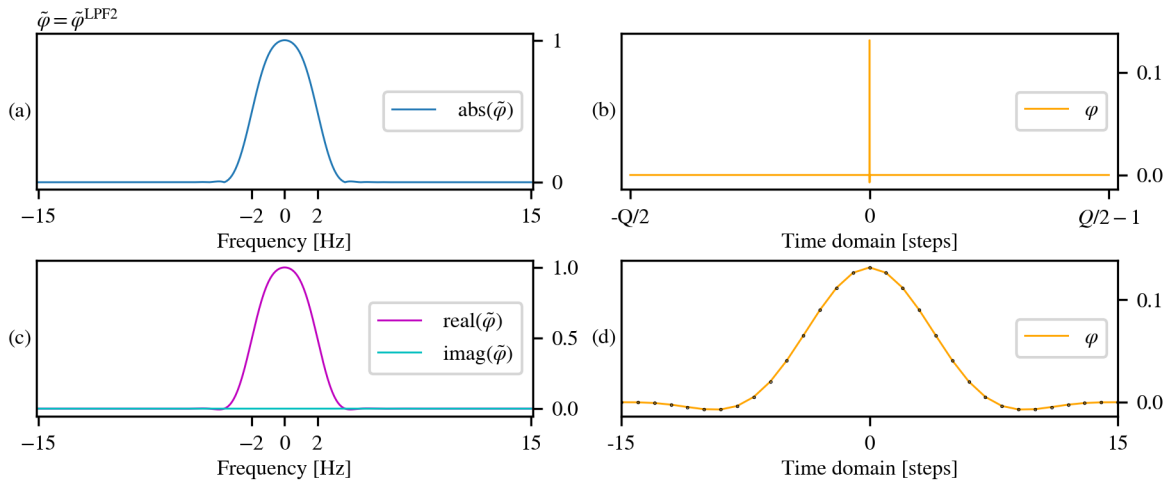

Figure 4.20: (a) Absolute value and (c) real and imaginary parts of filter vector $\tilde{\varphi}^{\mathrm{LPF2}}$ ($M = 15$); (b) and (d) show the component-wise absolute value of the inverse Fourier transform of $\tilde{\varphi}^{\mathrm{LPF2}}$.

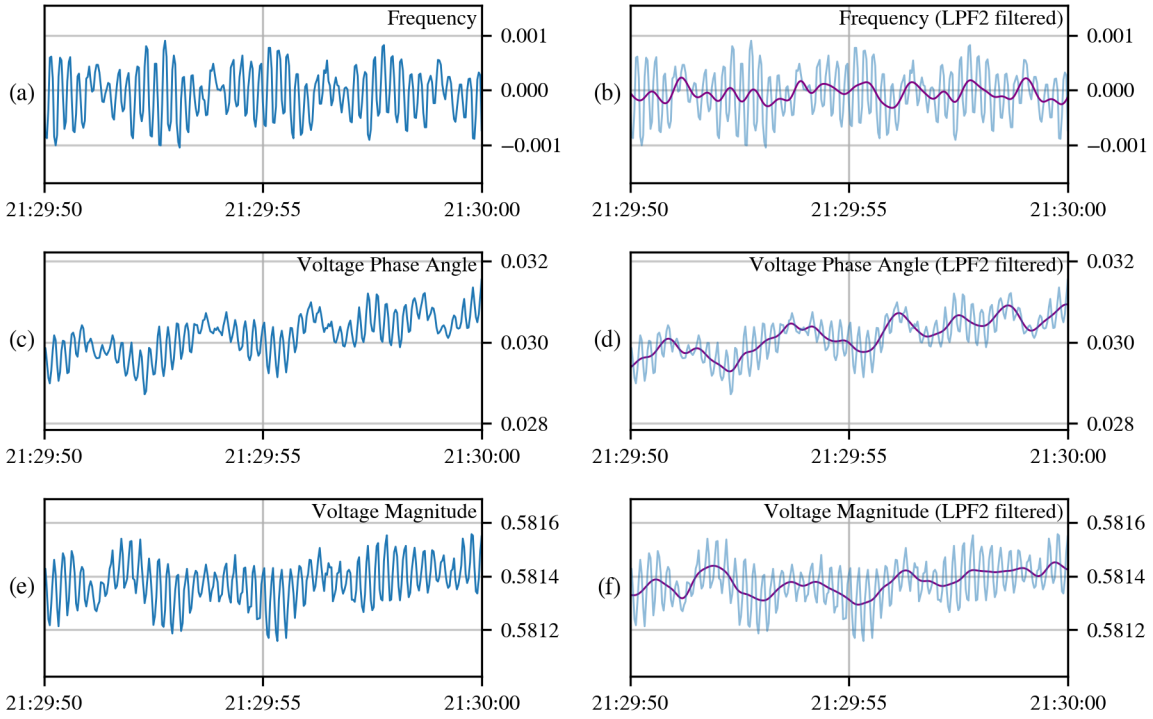

Figure 4.21: Original sampling (in blue) for (a) frequency, (c) voltage phase angle, and (e) voltage magnitude across 10 seconds on period #10 (see Table 4.1). Corresponding filtered series are shown (in purple) in plots (b), (d), and (f) using filtering vector $\tilde{\varphi}^{\mathrm{LPF2}}$.

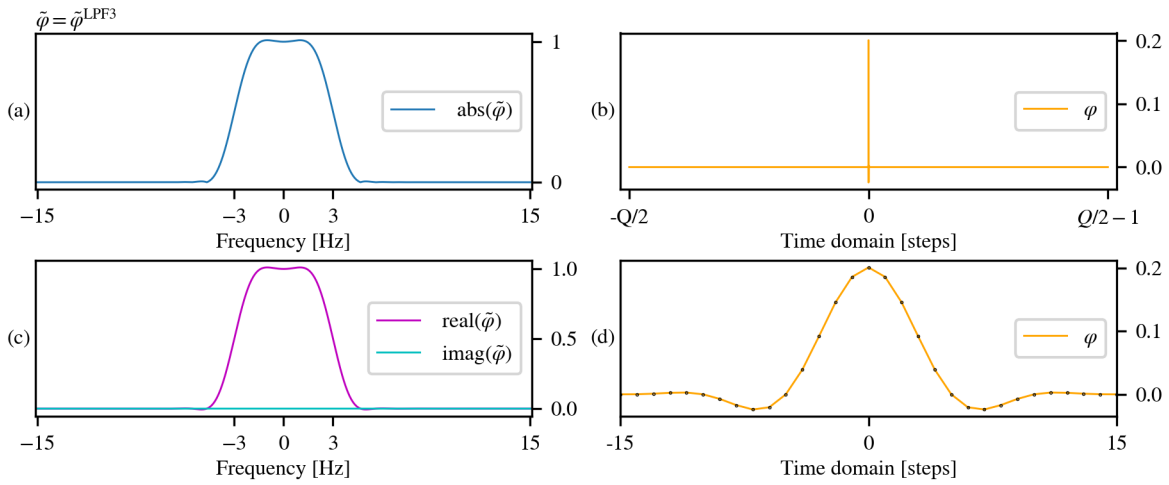

Figure 4.22: (a) Absolute value and (c) real and imaginary parts of filter vector $\tilde{\varphi}^{\mathrm{LPF3}}$ ($M = 15$); (b) and (d) show the component-wise absolute value of the inverse Fourier transform of $\tilde{\varphi}^{\mathrm{LPF3}}$.

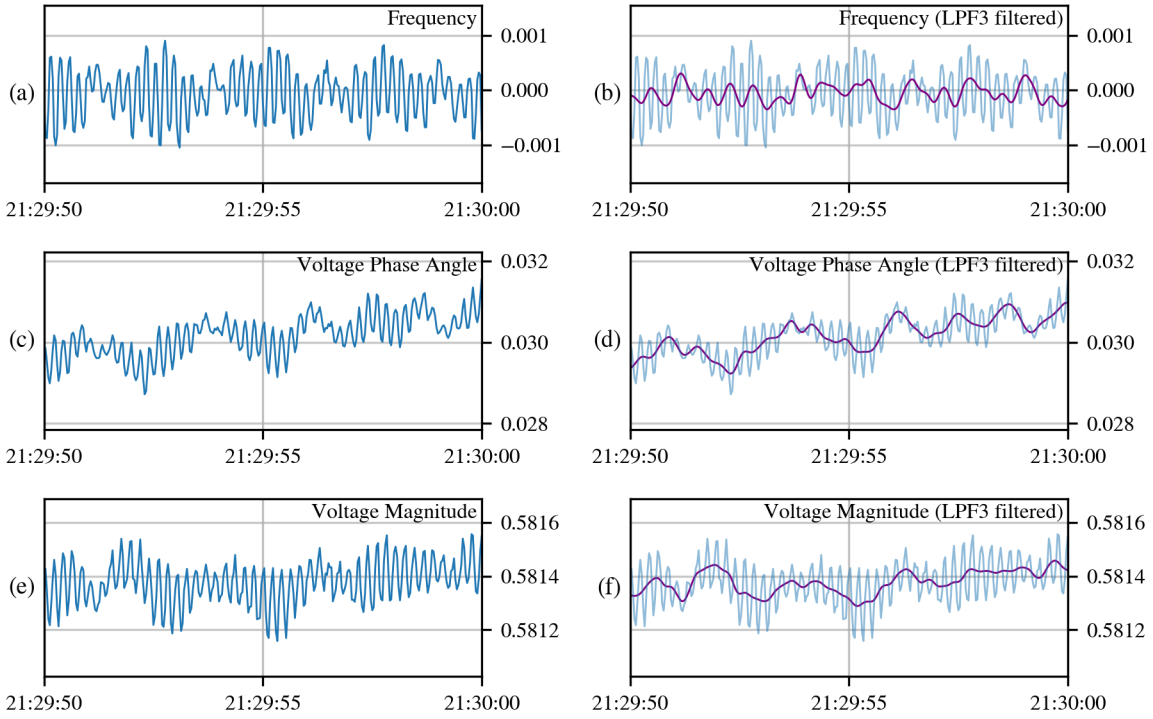

Figure 4.23: Original sampling (in blue) for (a) frequency, (c) voltage phase angle, and (e) voltage magnitude across 10 seconds on period #10 (see Table 4.1). Corresponding filtered series are shown (in purple) in plots (b), (d), and (f) using filtering vector $\tilde{\varphi}^{\mathrm{LPF3}}$.

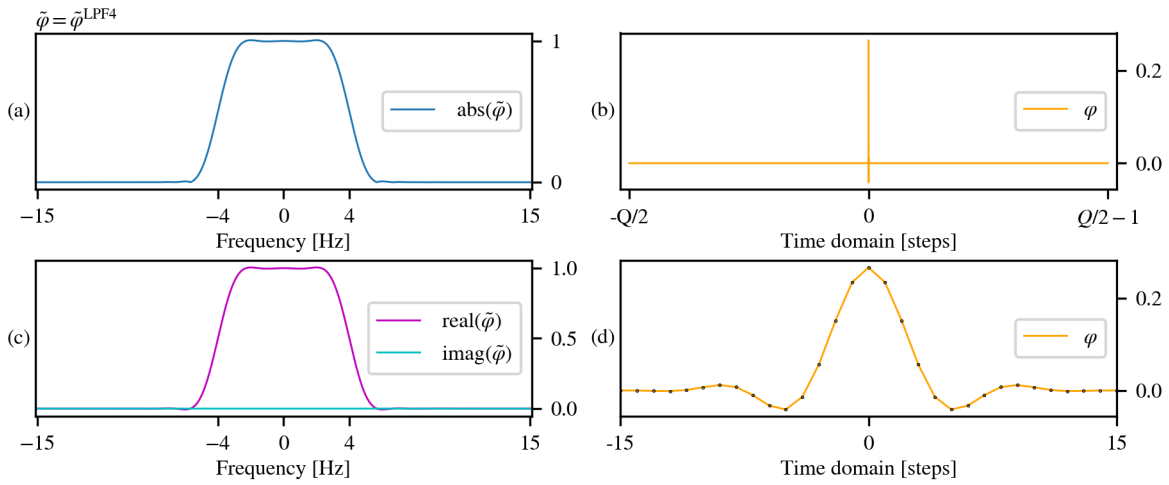

Figure 4.24: (a) Absolute value and (c) real and imaginary parts of filter vector $\tilde{\varphi}^{\mathrm{LPF4}}$ ($M = 15$); (b) and (d) show the component-wise absolute value of the inverse Fourier transform of $\tilde{\varphi}^{\mathrm{LPF4}}$.

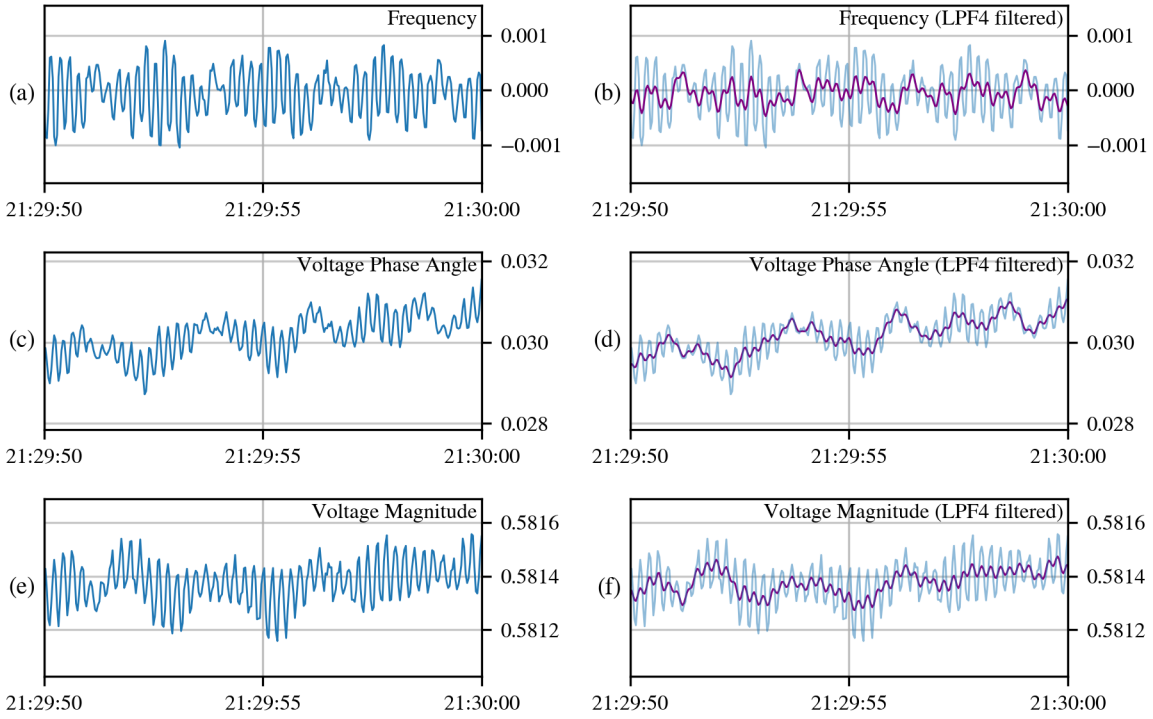

Figure 4.25: Original sampling (in blue) for (a) frequency, (c) voltage phase angle, and (e) voltage magnitude across 10 seconds on period #10 (see Table 4.1). Corresponding filtered series are shown (in purple) in plots (b), (d), and (f) using filtering vector $\tilde{\varphi}^{\mathrm{LPF4}}$.

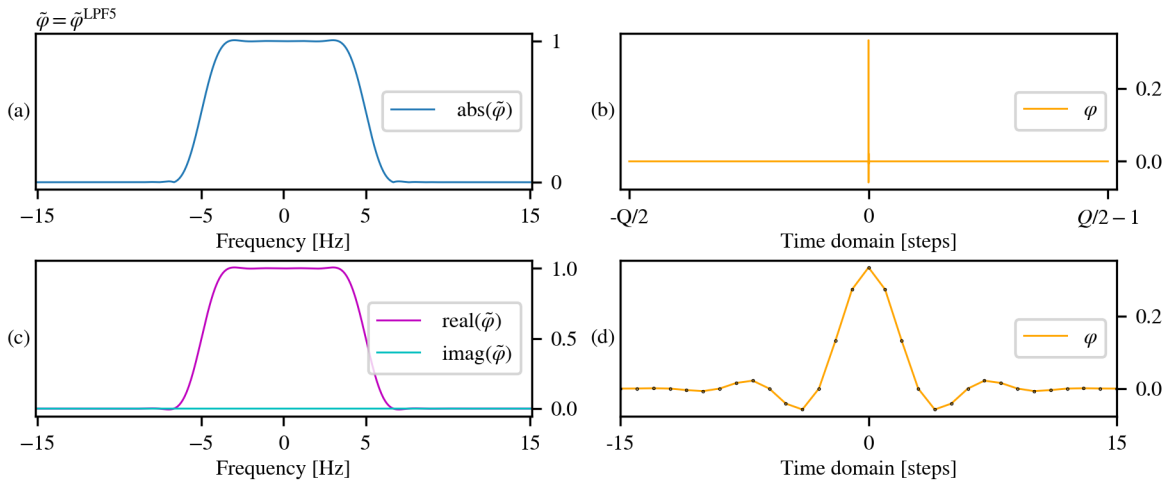

Figure 4.26: (a) Absolute value and (c) real and imaginary parts of filter vector $\tilde{\varphi}^{\mathrm{LPF5}}$ ($M = 15$); (b) and (d) show the component-wise absolute value of the inverse Fourier transform of $\tilde{\varphi}^{\mathrm{LPF5}}$.

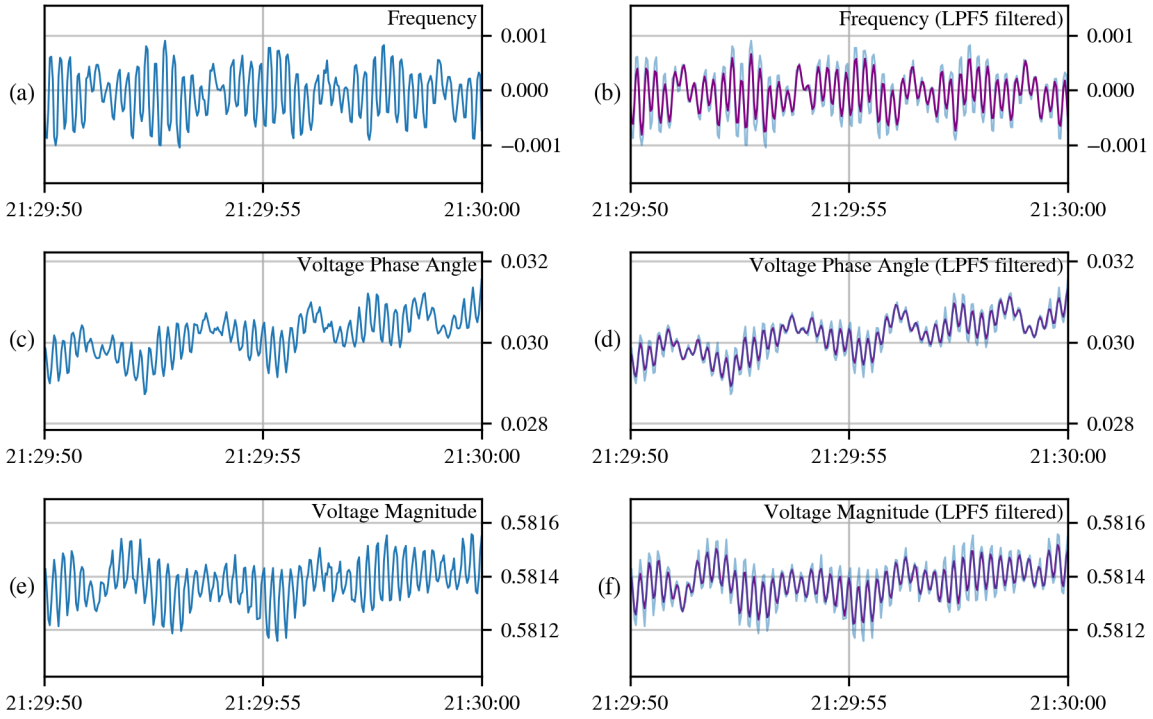

Figure 4.27: Original sampling (in blue) for (a) frequency, (c) voltage phase angle, and (e) voltage magnitude across 10 seconds on period #10 (see Table 4.1). Corresponding filtered series are shown (in purple) in plots (b), (d), and (f) using filtering vector $\tilde{\varphi}^{\mathrm{LPF5}}$.

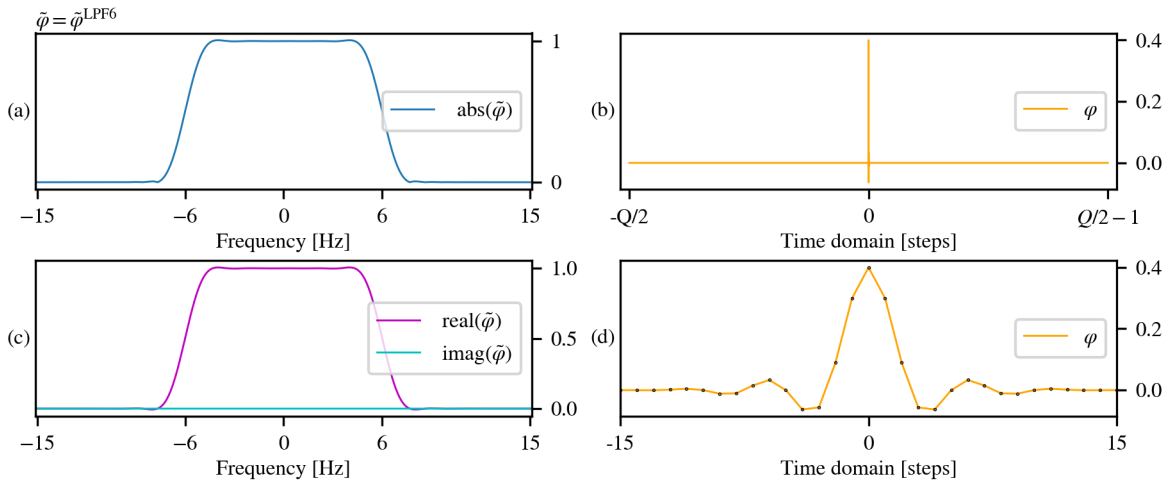

Figure 4.28: (a) Absolute value and (c) real and imaginary parts of filter vector $\tilde{\varphi}^{\mathrm{LPF6}}$ ($M = 15$); (b) and (d) show the component-wise absolute value of the inverse Fourier transform of $\tilde{\varphi}^{\mathrm{LPF6}}$.

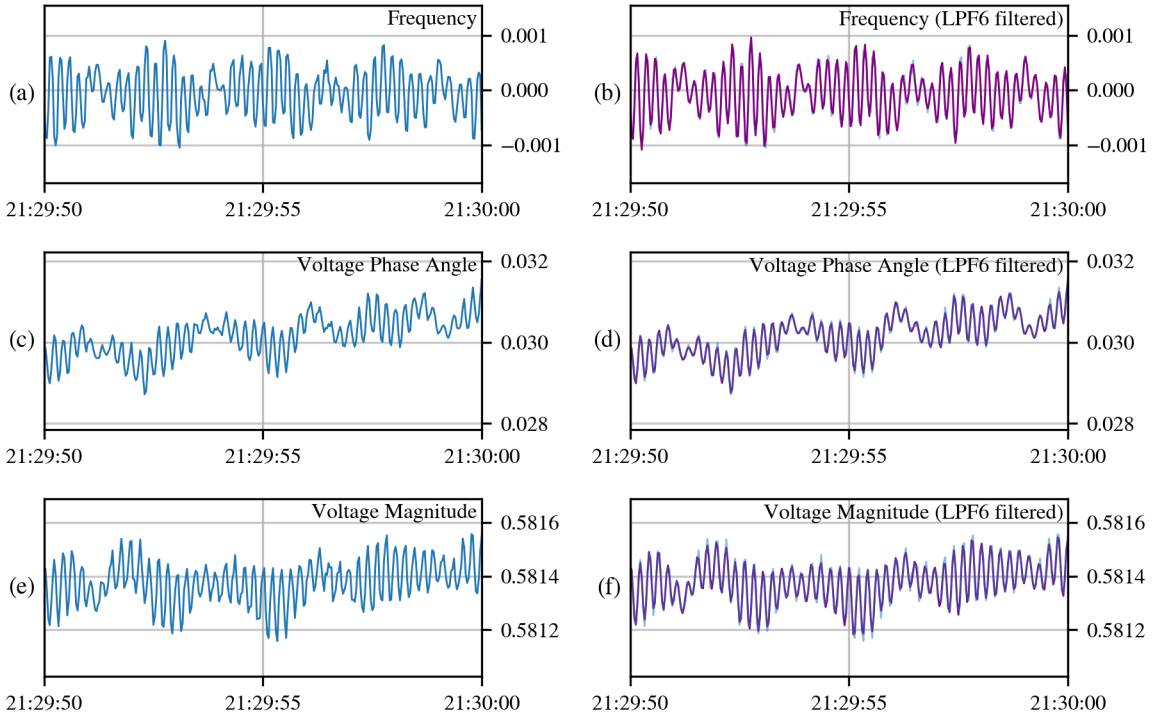

Figure 4.29: Original sampling (in blue) for (a) frequency, (c) voltage phase angle, and (e) voltage magnitude across 10 seconds on period #10 (see Table 4.1). Corresponding filtered series are shown (in purple) in plots (b), (d), and (f) using filtering vector $\tilde{\varphi}^{\mathrm{LPF6}}$.

We observe that there is prominent difference is the filtered series when frequencies larger than 5 Hz are erased from the original sampling. We will consider other two filter vectors to study and visualize the effect of the peaks near 5 Hz in the frequency domain of the original signal.

4.4.2 High Frequency Filter

Contrary to LPF we can also describe filter that keep the frequencies in the high range, these filters are called high-pass filters (HPF). Figure 4.30 describe the filter $\tilde{\varphi}^{\text{HPF5}}$ that cuts frequencies lower that 5 Hz, the filter is built given the $M = 15$ convolution coefficients depicted in plot (d). Figure 4.31 shows the resulting filtered series, note that now the filtered series have mean close to 0 (since the mean of the time series is represented by the zero-frequency coefficient of the Fourier transform, which is erased). Also note that the plots of the filtered series (on the right) have the same range length as the original time series plots (on the left). The oscillation showed on the filtered series might be obtained by the residue from the peaks around 5 Hz in the frequency domain of the original sampling, since the filter vector has mild slope.

4.4.3 Band-pass and Band-stop Filters

The two last filter that we will describe are called band-pass filter (BPF) and band-stop filter (BSF). The first one keeps the frequencies of the original time series (in the frequency domain) belonging to a particular band around a selected frequency; while the latter does the opposite, it suppresses the frequencies in that band. We will use these filters to specifically observe the contribution of the frequencies around the observed peaks at 5 Hz using $\tilde{\varphi}^{\text{BPF5}}$ (see Figures 4.32 and 4.33), and we remove them with $\tilde{\varphi}^{\text{BSF5}}$ (see Figures 4.34 and 4.35). In both cases, the filter vector have been described using $M = 30$ coefficients.

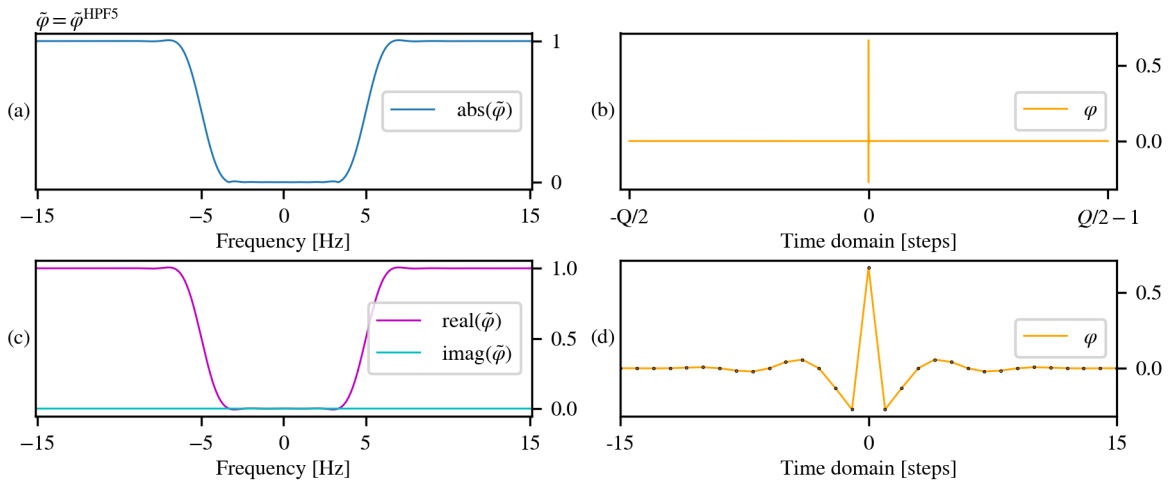

Figure 4.30: (a) Absolute value and (c) real and imaginary parts of filter vector $\tilde{\varphi}^{\mathrm{HPF5}}$ ($M = 15$); (b) and (d) show the component-wise absolute value of the inverse Fourier transform of $\tilde{\varphi}^{\mathrm{HPF5}}$.

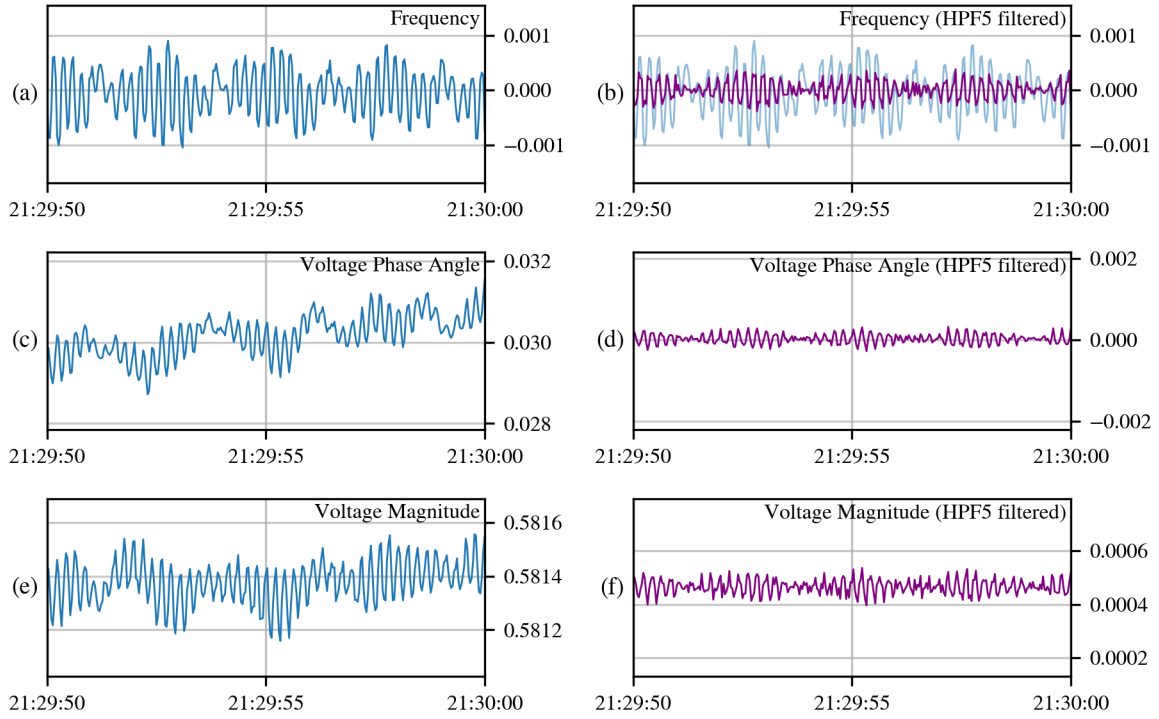

Figure 4.31: Original sampling (in blue) for (a) frequency, (c) voltage phase angle, and (e) voltage magnitude across 10 seconds on period #10 (see Table 4.1). Corresponding filtered series are shown (in purple) in plots (b), (d), and (f) using filtering vector $\tilde{\varphi}^{\mathrm{HPF5}}$.

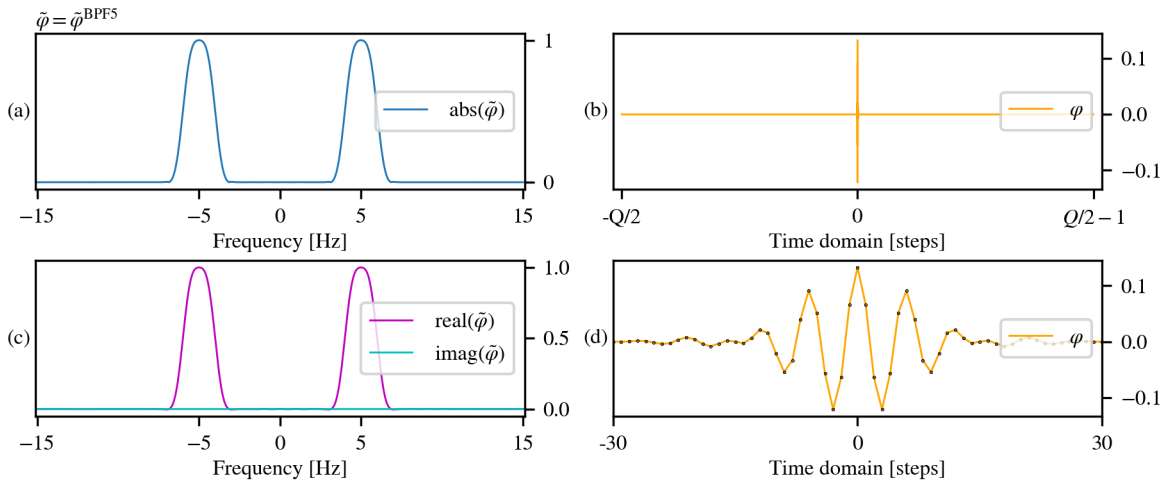

Figure 4.32: (a) Absolute value and (c) real and imaginary parts of filter vector $\tilde{\varphi}^{\mathrm{BPF5}}$ ($M = 30$); (b) and (d) show the component-wise absolute value of the inverse Fourier transform of $\tilde{\varphi}^{\mathrm{BPF5}}$.

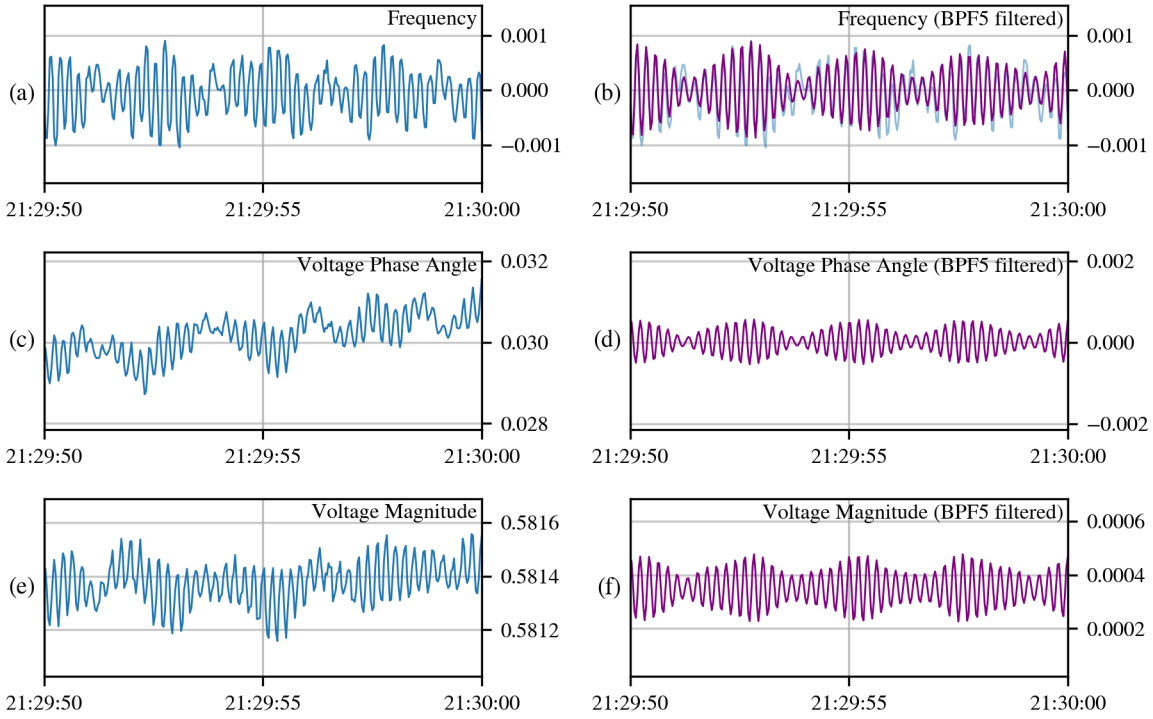

Figure 4.33: Original sampling (in blue) for (a) frequency, (c) voltage phase angle, and (e) voltage magnitude across 10 seconds on period #10 (see Table 4.1). Corresponding filtered series are shown (in purple) in plots (b), (d), and (f) using filtering vector $\tilde{\varphi}^{\mathrm{BPF5}}$.

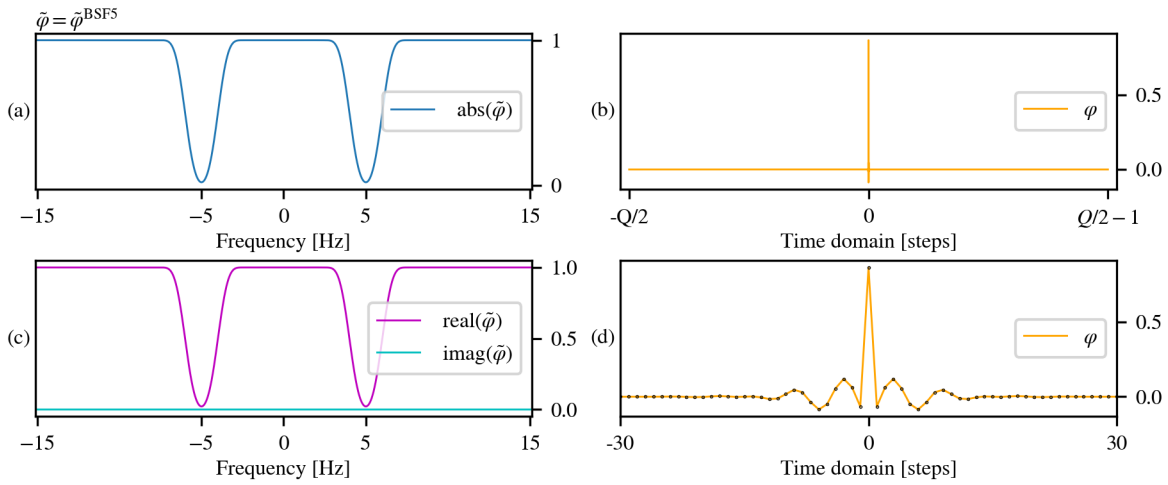

Figure 4.34: (a) Absolute value and (c) real and imaginary parts of filter vector $\tilde{\varphi}^{\mathrm{BSF5}}$ ($M = 30$); (b) and (d) show the component-wise absolute value of the inverse Fourier transform of $\tilde{\varphi}^{\mathrm{BSF5}}$.

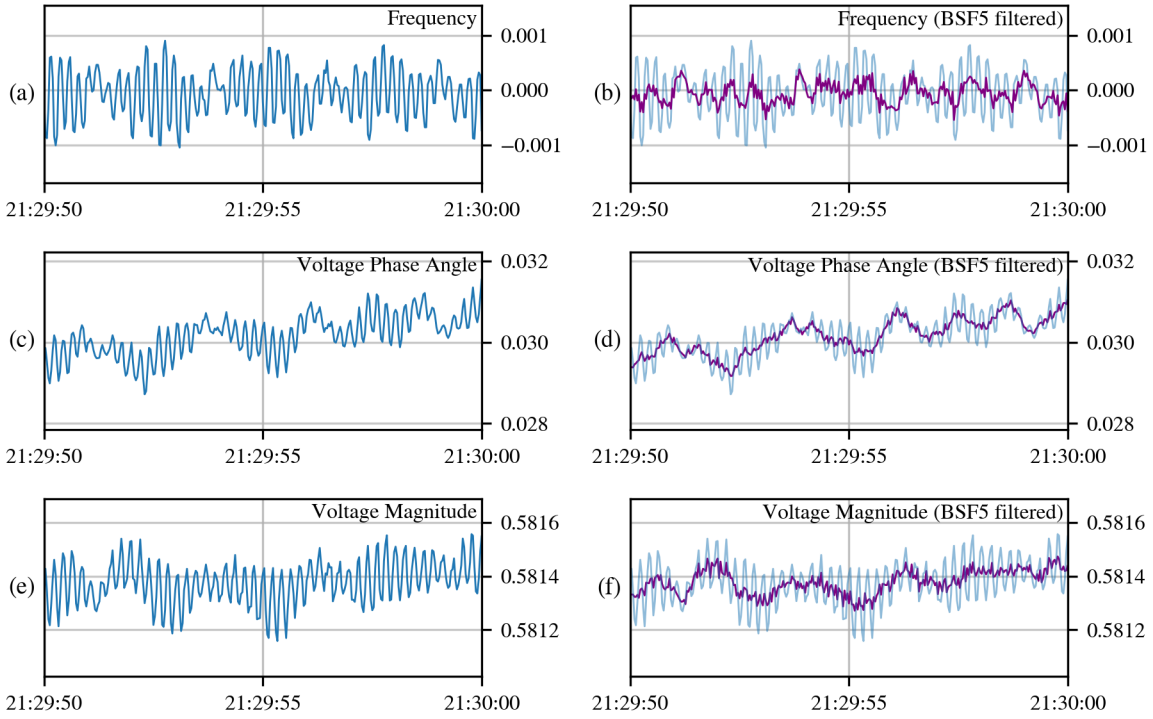

Figure 4.35: Original sampling (in blue) for (a) frequency, (c) voltage phase angle, and (e) voltage magnitude across 10 seconds on period #10 (see Table 4.1). Corresponding filtered series are shown (in purple) in plots (b), (d), and (f) using filtering vector $\tilde{\varphi}^{\mathrm{BSF5}}$.

4.5 Covariance Matrix and Principal Component Analysis

In this section we will describe a method to study the correlation of the measurements. We start by defining a covariance matrix of the sampling, which then is analyzed in terms of its spectrum (eigen-values and eigen-vectors), specifically we will observe how these parameters change over time.

Let $T \leq Q$ be the number of measurements of a time period. (Recall that Q is the number of measurements in a quiet period, that in our analyses has been set to be of 15-minute long.) The $N \times N$ covariance (correlation) matrix of the signal, which is based on the last T measurements, is defined by:

$$\begin{aligned} &\forall \text{ time } t \in \{T, \ldots, Q\}, \\ &\qquad \Sigma_0(t; T; m^*(\cdot)) \;\doteq\; \Big[\; \frac{1}{T} \sum_{\tau=t-T+1}^{t} m_k^*(\tau)\overline{m_\ell^*(\tau)} \;\Big|\; \forall k, \ell \in \{1, \ldots, N\} \; \Big] \\ &\qquad\qquad = \; \frac{1}{T} \sum_{\tau=t-T+1}^{t} m^*(\tau) m^*(\tau)^{\mathsf{H}}, \end{aligned} \tag{4.19}$$

where $m^*(\cdot)$ could be replaced by any of the normalized series defined above, that is, $\bar{m}^{(m)}(\cdot;\alpha)$, $\hat{m}^{(m)}(\cdot;\alpha)$, $\bar{m}^{(s)}(\cdot;S)$, $\hat{m}^{(s)}(\cdot;S)$, $\bar{f}^{(s)}[\tilde{\varphi}](\cdot;S)$, or $\hat{f}^{(s)}[\tilde{\varphi}](\cdot;S)$. It is easy to check that $\Sigma_0(t;T;m^*(\cdot))$ is a positive semidefinite matrix: let $z \in \mathbb{C}^N$, then

$$z^{\mathsf{H}}\, \Sigma_0(t;T;m^*(\cdot))\, z \;=\; \frac{1}{T} \sum_{\tau=t-T+1}^{t} z^{\mathsf{H}} m^*(\tau) m^*(\tau)^{\mathsf{H}} z \;=\; \frac{1}{T} \sum_{\tau=t-T+1}^{t} \left| z^{\mathsf{H}} m^*(\tau) \right|^2 \;\geq\; 0.$$

We aim to choose a (fixed) value T that is sufficiently large so that the covariance is weakly dependent on T, but also not too large to keep memory as small as possible (with an eye on streaming applications). Empirical experiments show that with T corresponding to the number of measurements in three minutes we can obtain stable covariance matrices when t varies.

In what follows we will drop the inputs T and $m^*(\cdot)$ from $\Sigma_0(t;T;m^*(\cdot))$ since they become fixed in the following analysis. We perform an eigen-decomposition on the cor-

relation matrix $\Sigma_0(t)$ and track the results as the function of t:

$$\Sigma_0(t) \;=\; \sum_{p=1}^{N} \lambda_p(t)\xi_p(t)\xi_p(t)^{\mathsf{H}}, \tag{4.20}$$

where $\xi_p(t)$ and $\lambda_p(t)$ are the orthonormal eigenvectors and corresponding eigenvalues (in decreasing order) of $\Sigma_0(t)$ respectively, i.e.

$$\forall p \in \{1,\dots,N\}: \quad \Sigma_0(t)\xi_p(t) \;=\; \lambda_p(t)\xi_p(t), \tag{4.21}$$

$$\forall p,q \in \{1,\dots,N\}: \quad \xi_p(t)^{\mathsf{H}}\xi_q(t) \;=\; \delta_{pq}, \tag{4.22}$$

$$\lambda_1(t) \geq \lambda_2(t) \geq \cdots \geq \lambda_N(t) \geq 0, \tag{4.23}$$

where δ is the Kronecker delta function ($\delta_{pq} = 1$ if and only if $p = q$). Since $\Sigma_0(t)$ is a positive semidefinite matrix, (4.23) is justified. In our tests, $\Sigma_0(t)$ is (numerically) rank-deficient, thus justifying the Principal Component Analysis (PCA) approach. PCA may be considered exact, if $P < N$ principal components are tracked or approximated.

The results of the PCA analysis are presented in Figures 4.36–4.41 which are screenshots of the movies available in [20]. The screenshots and the movies show at time t for each of the three variables (frequency, voltage angle, and voltage magnitude) four indicators:

- normalized vector $m^*(t)$ of measurements: for each sensor, we plot at its geographical position the value of the normalized data using different colors for different values;
- first 40 eigenvalues of $\Sigma_0(t)$, in decreasing order;
- largest 4 eigenvalues $\lambda_1(\cdot), \lambda_2(\cdot), \lambda_3(\cdot), \lambda_4(\cdot)$ for the last minute before t;
- corresponding eigenvectors $\xi_1(t), \xi_2(t), \xi_3(t), \xi_4(t)$ at time t, component values of the vector are plotted geographically, in blue values close to -1 and in red values close to 1.

For the numerical experiments we used traditional methods to compute the eigen-decomposition of $\Sigma_0(t)$, typically taking $O(N^3)$ time and using $O(NT)$ space. However, for streaming data that falls into a subspace of the original space (like in our case), we can consider lighter and faster algorithms to compute PCA, see [22].

From the simulations we can conclude that with few eigen-vector (say, less than 10) we can characterize more than 80% of the spectrum of the correlation matrix. Moreover, the largest eigen-values and the corresponding eigen-vectors are stable over time during the quiet periods. All these characteristics (correlation matrix, eigen-values and eigen-vectors) show spatial relationship between the measurements (that is, the relation that exists between different PMU sensors for a period of time). In what follows, we argue that we can also obtain temporal correlation across sensors.

4.5.1 Singular Value Decomposition

Consider the $N \times T$ measurement matrix

$$M(t) \;=\; [\; M_{k,\tau} = m_k^*(\tau) \mid \forall k \in \{1,\ldots,N\},\; \forall \tau \in \{t-T+1,\ldots,t\}\;]\,.$$

The singular value decomposition (SVD) of the measurement matrix $M(t)$ is

$$M(t) \;=\; U(t)D(t)W(t)^{\mathsf{H}}, \tag{4.24}$$

where the "spatial" matrix $U(t)$ is an $N \times N$ matrix the columns of which are orthogonal unit vectors of length N which are called left singular vectors of $M(t)$, the "temporal" matrix, $W(t)$, is $S \times S$ whose columns are right singular vectors of $M(t)$ and $D(t)$ is the $N \times S$ rectangular diagonal matrix of positive numbers. The covariance matrix, $\Sigma_0(t)$, is related to the measurement matrix, $M(t)$, according to

$$\Sigma_0(t) \;=\; M(t)M(t)^{\mathsf{H}} \;=\; U(t)D(t)D(t)^{\mathsf{H}}U(t)^{\mathsf{H}}, \tag{4.25}$$

where we took into account that $W(t)^{\mathsf{H}}W(t) = I$. One observes that the covariance matrix does not depend on the temporal matrix $W(t)$. To investigate time-related correlation effects, in the next section we also study auto-correlations.

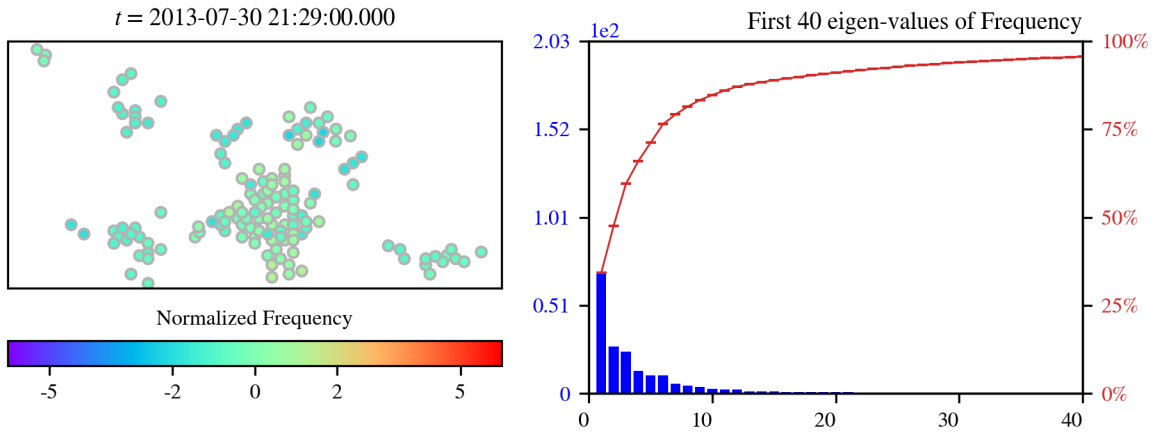

Figure 4.36: Left: normalized frequency $\hat{m}^{(s)}(t;S)$ representing each PMU in its geographical location. Right: First 40 eigen-values of $\Sigma_0(t;T;\hat{m}^{(s)}(\cdot;S))$ in blue, and its spectral contribution in red. [t=21:29:00 on July 30, 2013; $S=30$; $T=5400$]

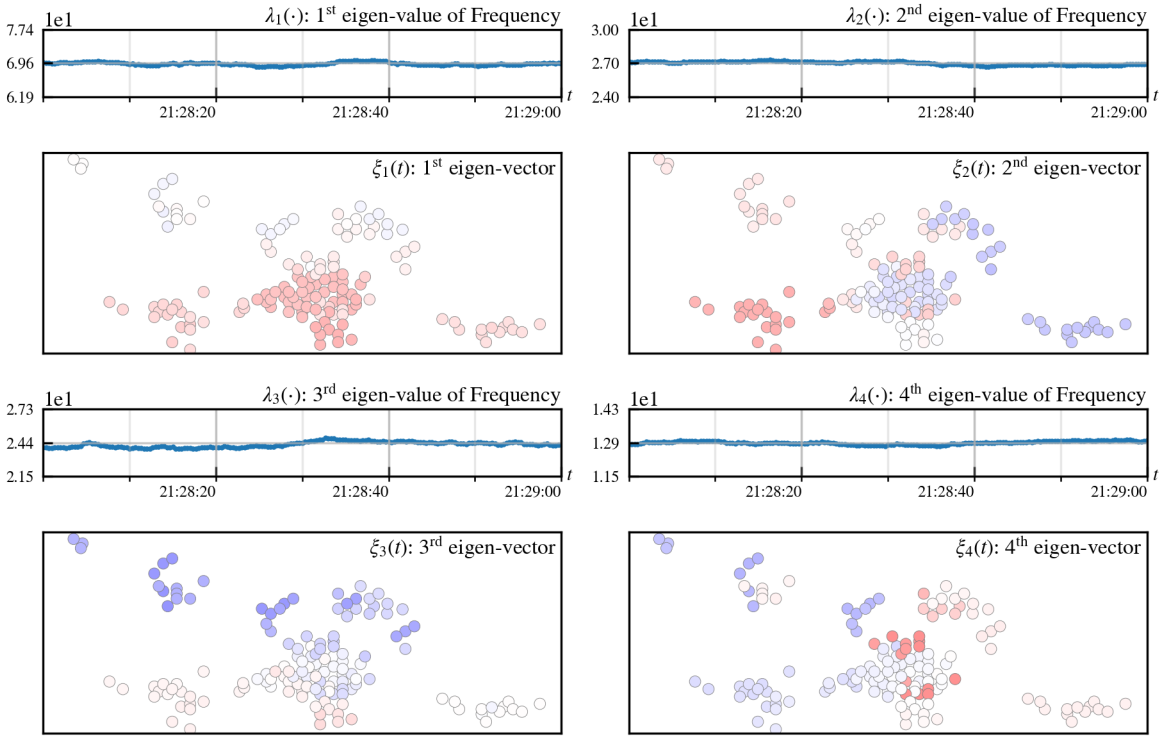

Figure 4.37: Largest four eigen-values from Figure 4.36 across the last minute before t, together with their corresponding normalized eigen-vectors at time t (positive values are shown in red, negative values are shown in blue). [t=21:29:00 on July 30, 2013]

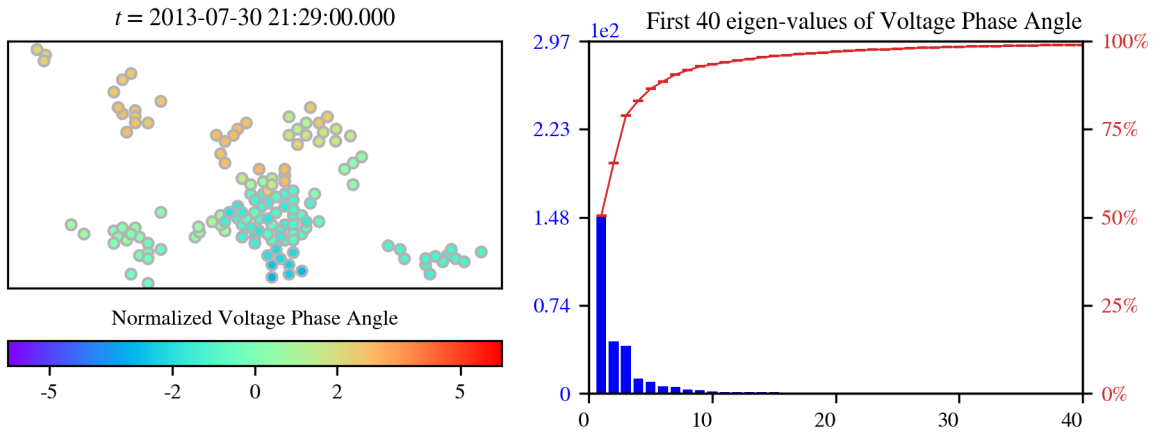

Figure 4.38: Left: normalized voltage phase angle $\hat{m}^{(s)}(t;S)$ representing each PMU in its geographical location. Right: First 40 eigen-values of $\Sigma_0(t;T;\hat{m}^{(s)}(\cdot;S))$ in blue, and its spectral contribution in red. [t=21:29:00 on July 30, 2013; $S=30$; $T=5400$]

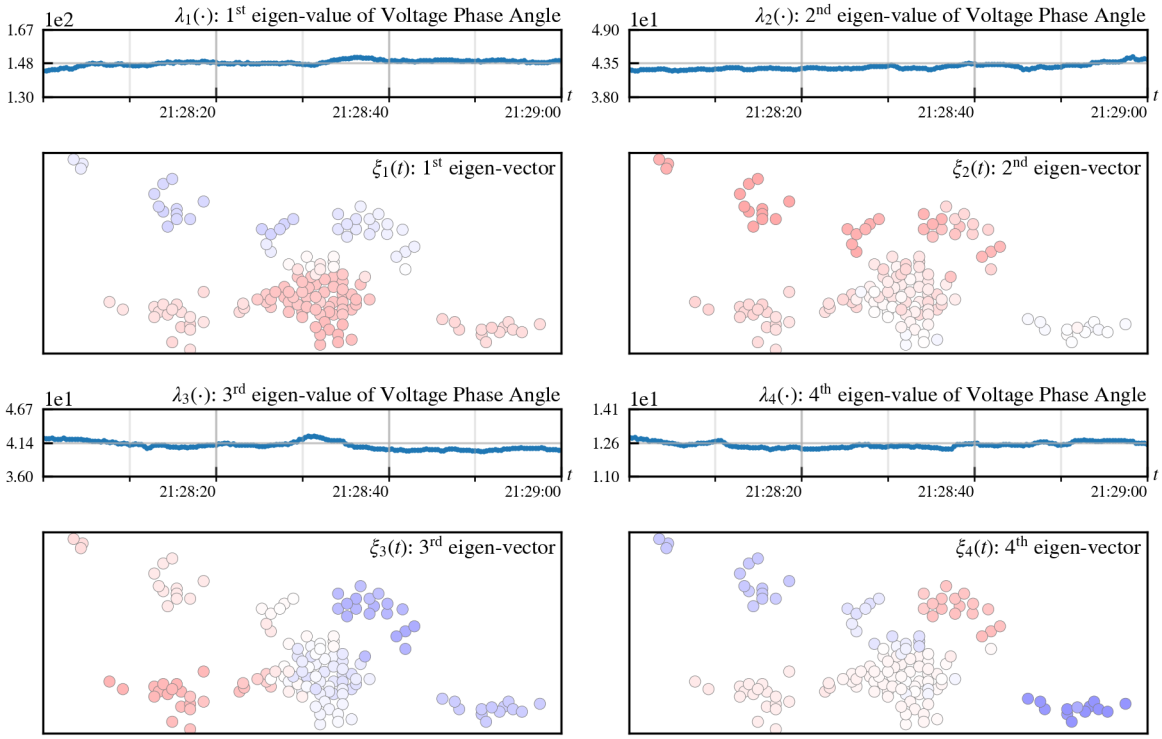

Figure 4.39: Largest four eigen-values from Figure 4.38 across the last minute before t, together with their corresponding normalized eigen-vectors at time t (positive values are shown in red, negative values are shown in blue). [t=21:29:00 on July 30, 2013]

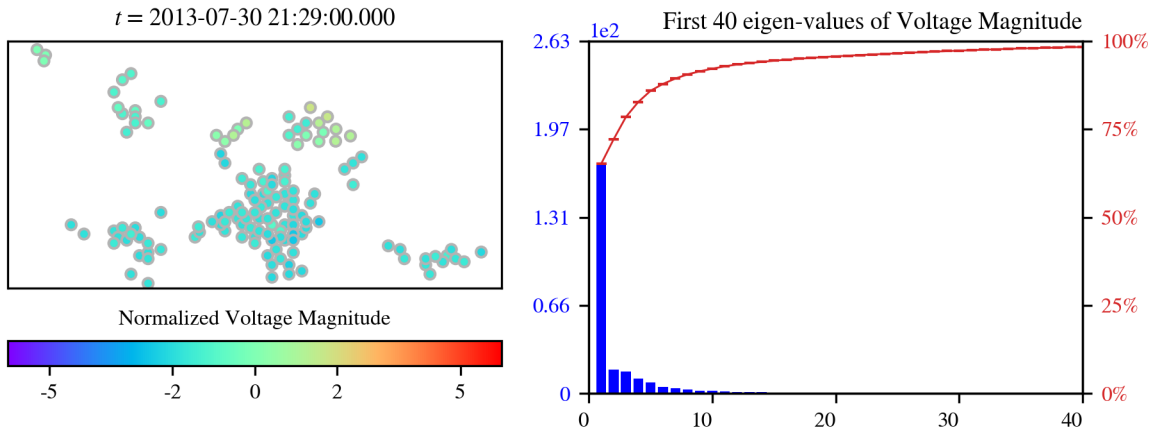

Figure 4.40: Left: normalized voltage magnitude $\hat{m}^{(s)}(t;S)$ representing each PMU in its geographical location. Right: First 40 eigen-values of $\Sigma_0(t;T;\hat{m}^{(s)}(\cdot;S))$ in blue, and its spectral contribution in red. [t=21:29:00 on July 30, 2013; $S=30$; $T=5400$]

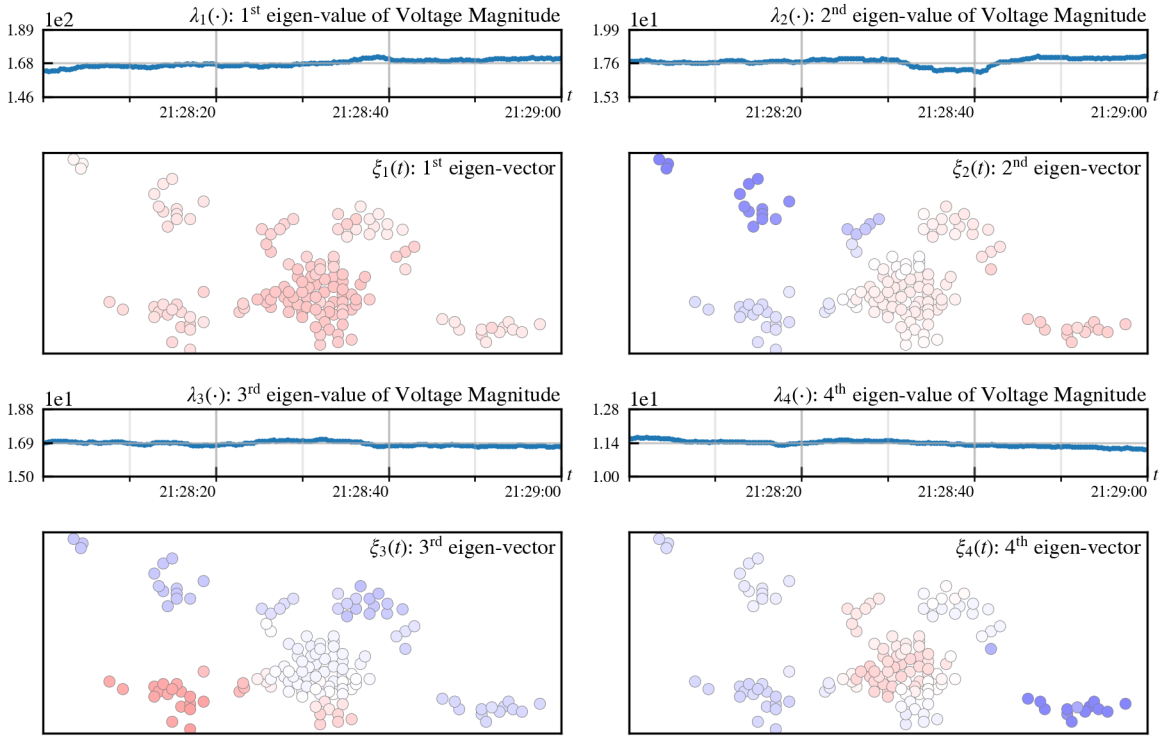

Figure 4.41: Largest four eigen-values from Figure 4.40 across the last minute before t, together with their corresponding normalized eigen-vectors at time t (positive values are shown in red, negative values are shown in blue). [t=21:29:00 on July 30, 2013]

4.6 Accounting for Temporal Correlations

In this section we describe a method to obtain temporal correlation between the PMU sensors.

Consider the delayed covariance matrix generalizing Eq. (4.19):

$$\forall \Delta \geq 0,\ \forall t \geq T + \Delta,$$

$$\begin{aligned} \Sigma_\Delta(t; T; m^*(\cdot)) &\doteq \left[\frac{1}{T} \sum_{\tau=t-T+1}^{t} m_k^*(\tau)\overline{m_\ell^*(\tau-\Delta)} \;\middle|\; \forall k, \ell \in \{1, \ldots, N\} \right] \\ &= \frac{1}{T} \sum_{\tau=t-T+1}^{t} m^*(\tau) m^*(\tau-\Delta)^{\mathsf{H}}. \end{aligned} \tag{4.26}$$

We are interested to study how $\Sigma_\Delta(t; T; m^*(\cdot))$ changes when Δ increases and then track evolution with t. Note that, in comparison with Σ_0, Σ_Δ is not necessarily positive semidefinite when $\Delta > 0$.

However, evolution of the spectrum (in the two-dimensional (Δ, t) space) is challenging. Instead, we study two surrogate objects, introduced in the following two subsections, which are easier to visualize. Again, since T and $m^*(\cdot)$ are fixed, we will omit them as input of the Σ_Δ function.

4.6.1 Auto-Correlation Functions

To study the correlation between measurements on a sensor and its own past we introduce the auto-correlation functions.

Definition 21. *The normalized PMU's* auto-correlation functions *at different nodes are defined as follows:*

$$\forall k \in \{1, \ldots, N\}, \qquad \mathcal{A}_k(\Delta; t) = \frac{[\Sigma_\Delta(t)]_{kk}}{[\Sigma_0(t)]_{kk}}. \tag{4.27}$$

We choose to normalize the function with respect to the term $[\Sigma_0(t)]_{kk}$ so we can have a relative comparison tool for different sensor. It might be useful to think about these

auto-correlation functions for a fixed t and see it as a function of Δ, that is $\mathcal{A}_k(\cdot;t)$. This object is of interest because of the following two reasons

- Dependence of the auto-correlation function on Δ indicates whether fluctuations around the mean at a particular node decay or not with time. Stated differently this analysis tests if there are significant memory effects or if memory is lost.
- It accounts for the part of the measurement matrix which is ignored in the PCA analysis, as discussed above –that is, it accounts for the temporal matrix, W.

We show the auto-correlation function for frequency at a fixed t for four selected PMU sensors in Figure 4.42. We obtain different patters and shapes of the resulting auto-correlation functions: some of them have large amplitude, other ones have sinusoidal shape.

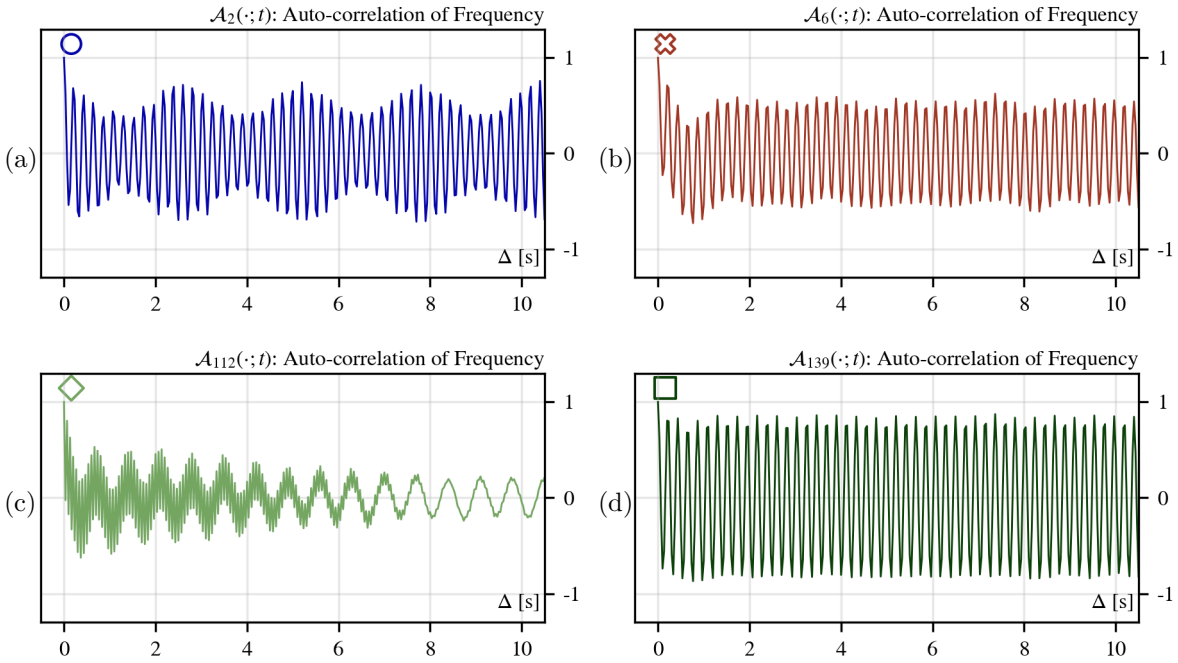

Figure 4.42: Auto-correlation functions for frequency for PMU's (a) $k = 2$, (b) $k = 6$, (c) $k = 112$, and (d) $k = 139$; at t=21:29:00 on July 30, 2013. The correlation matrices are constructed with the normalized time series $\hat{m}^{(s)}(\cdot;S)$, $S = 30$, $T = 5400$.

We will compare these functions with the one that we obtain if we first apply a band-stop and band-pass Fourier filters at 5 Hz to the original sampling, with the objective of suppressing the oscillations that we observe at that frequency, and then observe their contribution on the these auto-correlation functions.

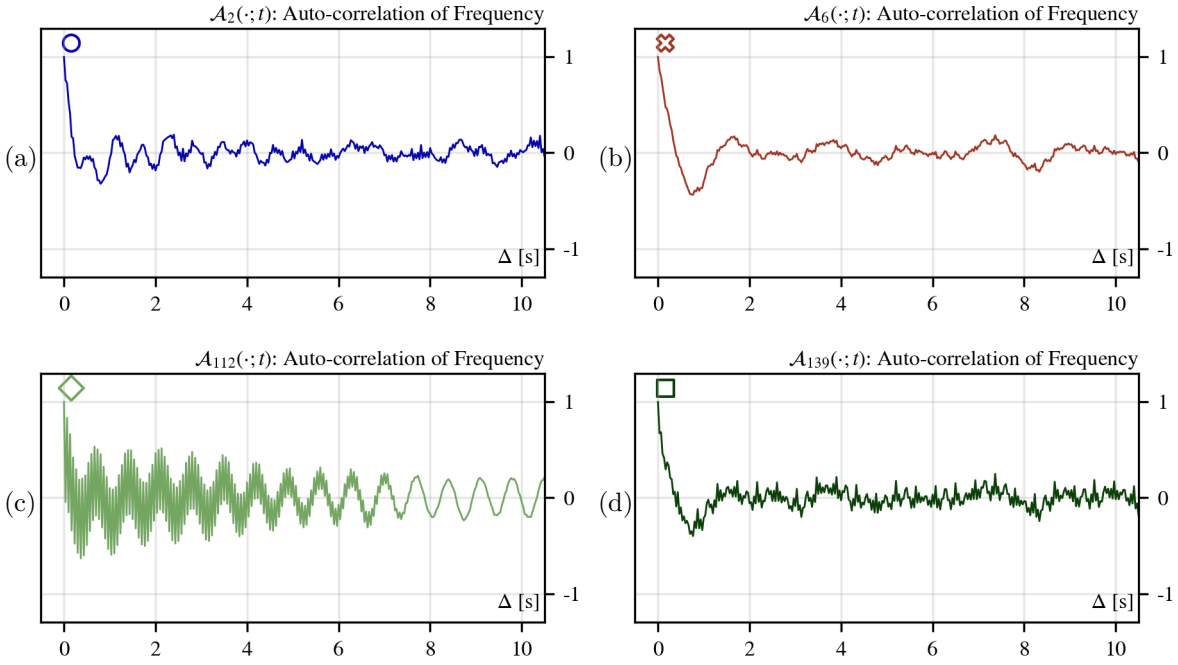

Figure 4.43: Auto-correlation functions for frequency for PMU's (a) $k = 2$, (b) $k = 6$, (c) $k = 112$, and (d) $k = 139$; at t=21:29:00 on July 30, 2013. The correlation matrices are constructed with the normalized time series $\hat{f}^{(s)}[\tilde{\varphi}^{\mathrm{BSF5}}](\cdot;S)$, $S = 30$, $T = 5400$.

Figure 4.43 shows the auto-correlation functions when a band-stop Fourier filter is applied (see details of the filter in Figure 4.34). We note that the oscillations and the amplitude of three out of the four cases that are depicted drop down, therefore, the reason of the high auto-correlation of the measurements might be explained by the 5 Hz oscillation of the sampling —as we have seen, shown as peaks in the frequency domain of the Fourier transform.

Figure 4.44 illustrates the auto-correlation functions when a band-pass Fourier filter is applied, the frequencies that are kept in the spectrum are shown in Figure 4.32. We

observe that for the three sensors where the band-stop filter made a difference, the band-stop filter keeps the same shape of the functions as the original series (without filtering) amplifying their amplitude; however, for sensor $k = 112$ on Figure 4.44c, the auto-correlation functions changes completely, implying that the cause of the auto-correlation is out of the 4-6 Hz range in the frequency domain of the signal.

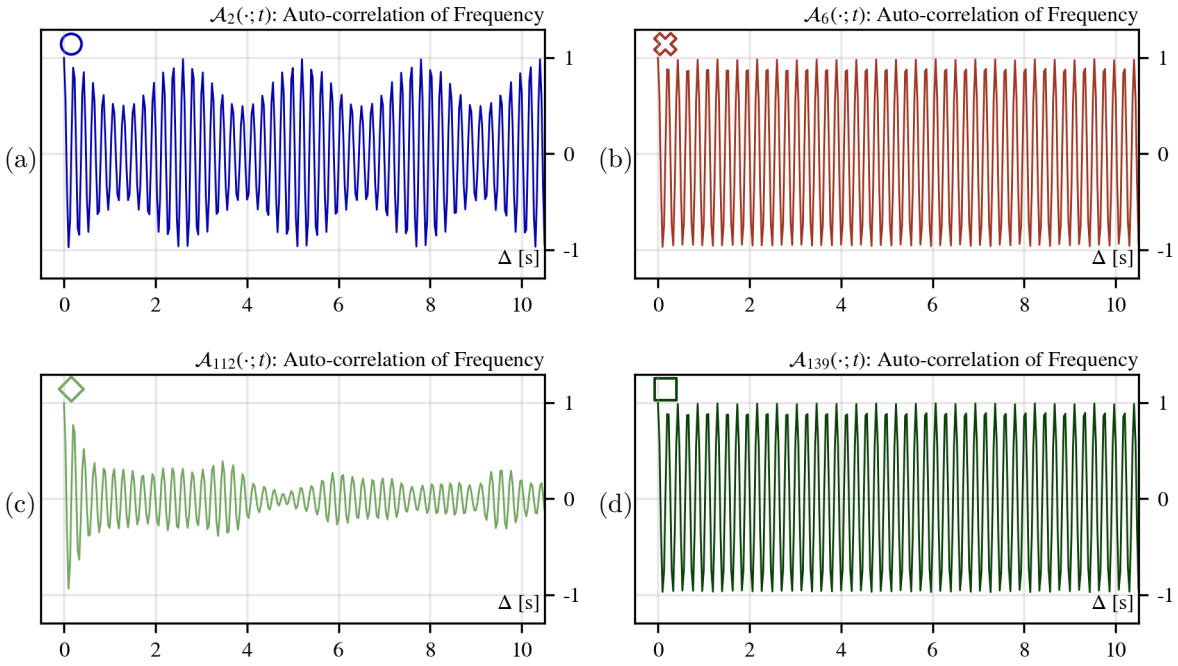

Figure 4.44: Auto-correlation functions for frequency for PMU's (a) $k = 2$, (b) $k = 6$, (c) $k = 112$, and (d) $k = 139$; at t=21:29:00 on July 30, 2013. The correlation matrices are constructed with the normalized time series $\hat{f}^{(s)}[\tilde{\varphi}^{\mathrm{BPF5}}](\cdot;S)$, $S = 30$, $T = 5400$.

We show study dependence of the auto-correlations on time, i.e. dependence of $\mathcal{A}_k(\cdot;t)$ on t. The movies are available in [20], showing examples of buses that have large auto-correlation amplitude. In general, the auto-correlation functions are stable on time (when t varies), indicating that the correlation of the sampling are consistent on quiet periods.

We are interested in understanding the fact that some buses have larger auto-correlations than others. This can be measured by the amplitude of the auto-correlation functions, we formalize this concept in the following definition:

Definition 22. *The* residue *of the auto-correlation functions is defined as*

$$\forall k \in \{1, \ldots, N\}, \quad \rho_k([\Delta_{\min}, \Delta_{\max}]; t) \doteq \frac{1}{\Delta_{\max} - \Delta_{\min}} \sum_{\Delta=\Delta_{\min}}^{\Delta_{\max}-1} |\mathcal{A}_k(\Delta; t)|, \qquad (4.28)$$

where $\Delta_{\min}$ *and* $\Delta_{\max}$ *are the initial and ending points of contribution to the residue.*

We plot the residue for each sensor in a geographical map, see Figure 4.45. In these computations, we have ignored the first second of the auto-correlation function —that represent the correlation of the sampling with its immediate past—, by setting $\Delta_{\min} = 30$. The maximum point of contribution has been set to be $\Delta_{\max} = 1800$ (1 minute), since we have empirically verified that with this value of $\Delta_{\max}$ the results that are obtained are robust with $O(1)$ changes). In the figure, we show the residual map obtained with the original normalized time series, using a band-stop filter, and using a band-pass filter.

As seen in the movies of the residue maps shown in Figure 4.45, residue values at a number of special nodes do not decay with time. Moreover we observe that nodes with a significant residue cluster, specifically, there are two sets of nodes that repeat their high values across different days, months and time of the day. Those groups are also shown in Figure 4.45a: in the bottom-left part of the map and the upper-right zone (closer to the center); they might show up together or by separate. We also note that the residue drastically drops when a band-stop filter is applied, as can be seen in Figure 4.45b; and it is accentuated when a band-pass filter is used, see Figure 4.45c.

Observation that sustainable correlations stay for sufficiently long period of time suggest to analyze spatio-temporal features of the sustainable correlations via the cross-correlation residue described in next.

Plots for the same auto-correlation analysis for the voltage phase angle and voltage magnitude measurements are shown in Figures A.1–A.8 in Appendices A.1 and A.2.

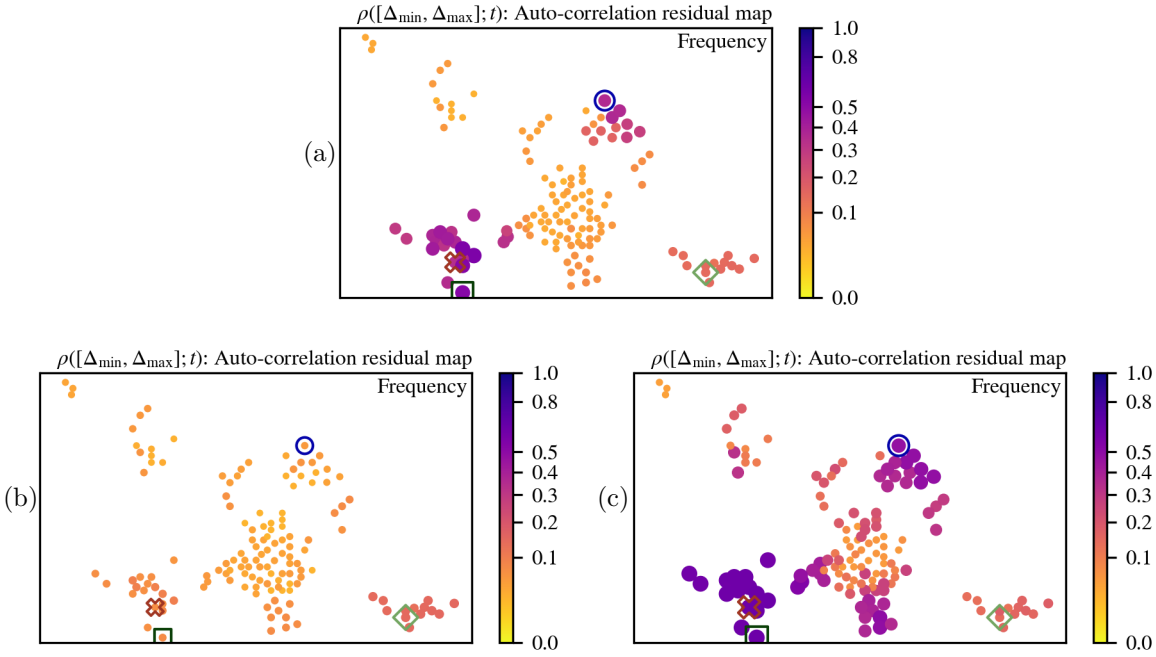

Figure 4.45: Residue of auto-correlation functions for frequency using (a) $\hat{m}^{(s)}(\cdot;S)$, (b) $\hat{f}^{(s)}[\tilde{\varphi}^{\mathrm{BSF5}}](\cdot;S)$, and (c) $\hat{f}^{(s)}[\tilde{\varphi}^{\mathrm{BPF5}}](\cdot;S)$; at t=21:29:00 on July 30, 2013, with $S = 30$, $T = 5400$, $\Delta_{\min} = 30$, and $\Delta_{\max} = 1800$. Geometrical figures show the position of the sensors depicted in Figures 4.42–4.44.

4.6.2 Cross-Correlation Residue (CCR)

The cross-correlation version of equations (4.27)–(4.28) is described in the following:

Definition 23. *The* cross-correlation functions *and their* residue *are defined, respectively, as*

$$\forall k, \ell \in \{1, \ldots, N\}, \qquad \mathcal{B}_{k\ell}(\Delta; t) \doteq [\Sigma_{\Delta}(t)]_{k\ell}, \tag{4.29}$$

$$\mathcal{R}_{k\ell}([\Delta_{\min}, \Delta_{\max}]; t) \doteq \frac{1}{\Delta_{\max} - \Delta_{\min}} \sum_{\Delta=\Delta_{\min}}^{\Delta_{\max}-1} |\mathcal{B}_{k\ell}(\Delta; t)|. \tag{4.30}$$

Here, we drop normalization in (4.29) to avoid singularities associated with signals at different nodes which are not correlated. To visualize the CCR (4.30), we plot the

cross-correlation residual matrix

$$\mathcal{R}([\Delta_{\min}, \Delta_{\max}]; t) \doteq \Big[\mathcal{R}_{k\ell}([\Delta_{\min}, \Delta_{\max}]; t) \mid \forall k, \ell \in \{1, \ldots, N\}\Big]. \tag{4.31}$$

In Figure 4.46 we plot the cross-correlation matrix (4.31) component-wise, using darker colors for higher values.

We observe that for the normalized original time series (see Figure 4.46a), there are some sensors that have high correlation with others, specifically, given the block structure of the matrix, there is a cluster of PMU sensors that have high correlation between each others. Figure 4.46b shows the cross-correlation matrix when a band-stop filter at 5 Hz is applied, dropping drastically the cross residue, which is consistent with what we observed with the auto-correlation. In Figure 4.46c we have applied a band-pass filter, and in this case, we see that the amount of sensors that are cross-correlated increases.

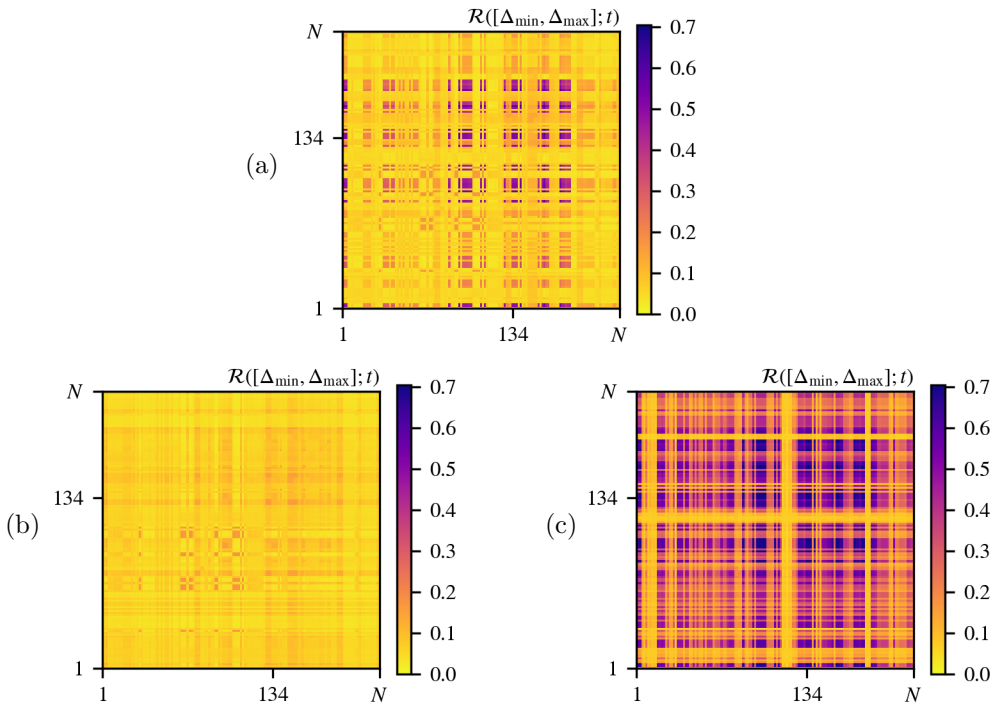

Figure 4.46: Cross-correlation residual matrix for frequency using (a) $\hat{m}^{(s)}(\cdot; S)$, (b) $\hat{f}^{(s)}[\tilde{\varphi}^{\mathrm{BSF5}}](\cdot; S)$, and (c) $\hat{f}^{(s)}[\tilde{\varphi}^{\mathrm{BPF5}}](\cdot; S)$; at t=21:29:00 on July 30, 2013, with $S = 30$, $T = 5400$, $\Delta_{\min} = 30$, and $\Delta_{\max} = 1800$.

For each of the three matrices plotted in Figure 4.46 we plot the components of the row corresponding to sensor $k = 134$ in order to see the where are the locations of the sensors that are cross-correlated with it. See Figure 4.47.

In the normalized original sampling, Figure 4.47a, we observe a high cross-correlation between the selected sensor and two clusters located at the bottom-left and upper-right of the map. Again, this structures repeats from the one observed in Figure 4.45a. As expected, no high cross-correlation is shown in the measurements filtered with a band-stop filter (in Figure 4.47b); however, we observe high cross-correlation between sensor $k = 134$ and most of the other sensor in Figure 4.47c. This might be related with the fact the the oscillations observed at 5 Hz are coordinated between most of the PMU's.

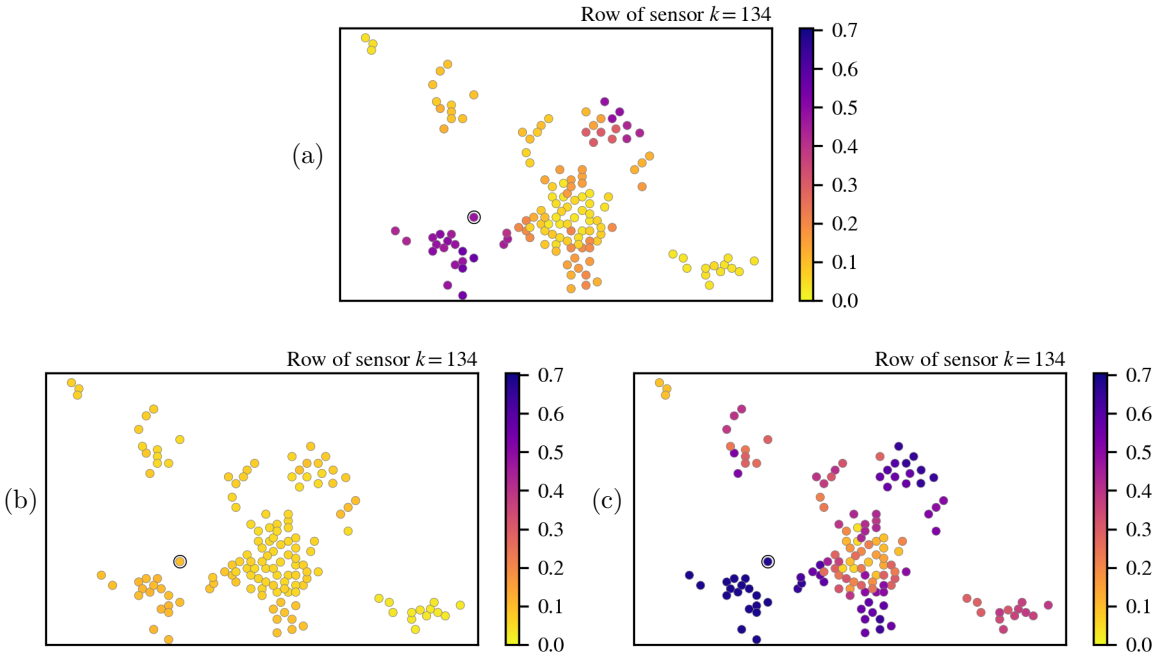

Figure 4.47: Cross-correlation residue for frequency, geographically located, between sensor $k = 134$ (indicated with an extra circle around it) and the remaining sensors using (a) $\hat{m}^{(s)}(\cdot; S)$, (b) $\hat{f}^{(s)}[\tilde{\varphi}^{\mathrm{BSF5}}](\cdot; S)$, and (c) $\hat{f}^{(s)}[\tilde{\varphi}^{\mathrm{BPF5}}](\cdot; S)$; at t=21:29:00 on July 30, 2013, with $S = 30$, $T = 5400$, $\Delta_{\min} = 30$, and $\Delta_{\max} = 1800$.

Plots for the same cross-correlation analysis for the voltage phase angle and voltage

magnitude measurements are shown in Figures A.9 and A.10 in Appendices A.3 and A.4.

4.7 Outcomes

We finalize this chapter by summarizing observations that we have made from the statistical analyzis that has been performed.

Averaged over time PMU signal shows interesting spatial correlations. Correlations are different for different objects of interest (frequency, phase and voltage) and also different for different quiet periods, altough general structures of the set of sensors that are correlated and their correlation functions are consistent over time. The matrix of correlations is sparse, also revealing that number of the high-intensity contributions to the correlations is small. (Note that the statement of sparseness is consistent with previous studies of the measurement matrix, see e.g. [57].) Each of the contributions characterizes a mode localized on a relatively few nodes (PMU positions) within the system. Principal modes, computed over the quiet periods, are almost frozen in time, however responding fast to any significant perturbation, thus suggesting them as efficient features/indicators for changes.

Fourier-analysis of the signal reveals interesting spatial patterns. Extracting modes in the 4-6 Hz range one observes significant contributions from only few PMU nodes. Whereas observing the contribution that the modes at the same range explains the high temporal auto- and cross-correlation of a significant portion of the sensors. Nodes showing large 4-6 Hz contributions were identified as aggregated loads, mid-size generators and large wind-farms.

In general the signals are long-correlated in time. However the memory effects becomes significantly less prominent when higher frequencies are removed. Nodes with significant auto-correlations (memory) have a spatial pattern is adiabatic/frozen (changing slowly during the quiet periods). Like other adiabatic patterns mentioned above, the pattern changes from one quiet period to another and we also observed slight different

patterns for different characteristics (frequency, phase, voltage).

Analyzing time-delayed cross-correlations between different nodes we observe that, like in the case of the auto-correlations, correlations between some nodes have long memory. Nodes which mutual inference shows a long memory form a sparse pattern. These patterns, like others described above are adiabatic and evolving from characteristic to characteristic and from one quiet period to another quiet period.

Chapter 5

Conclusion

During the first part of this thesis, we study cyber-physical attacks, attacks that are composed by a physical perturbation or disturbance in electrical grid simultaneously accompanied by a hack on the data produced by the sensors (PMUs) that measure the status of physical quantities of the network. Cyber-physical attacks have gained increased attention during the last decade since new technologies have become available and, therefore, their security systems are not strong enough yet. Hackers have discovered vulnerabilities and used them to harm the stability of the networks.

We introduce a sophisticated attack model, that includes load modifications and tripping of transmission lines in a small zone of the network, and avoids standard detection methods by injecting false measurements reflecting realistic safe flows. We are able to compute such attacks on large systems in seconds of CPU time, with hidden overloads of more than 50% of the lines capacity. If these attacks remain undetected for long periods of time, they could cause catastrophic consequences such as cascade line failures and extended blackouts.

We propose different stochastic defense mechanisms to augment standard detection tools for generic cyber attacks. Our defenses change the stochastics of system data using intelligent procedures that modify the power injection at specific generators in a way that is recognizable by the defender but difficult to anticipate by the attacker. The first

mechanism reveals the lines that connect the buses that are attacked with the rest of the network by using the change in generation to induce a change in the voltage phase angle of most of the nodes. Branches that show inconsistencies in terms of the voltage-current relationship are candidates for being at the boundary of the affected buses.

The second mechanism causes low-rank changes in the covariance (correlation) matrix of phase angles, affecting all the components of the matrix. The randomness during the iterations of the strategy and the complexity of the power flow equations makes an attacker's counteraction very difficult and, therefore, the nodes that were hacked will be exposed by not showing the proper low-rank covariance matrix correction.

The second half of this thesis is dedicated to the analysis of PMU measurements performed by an Independent System Operator in the United States. The available data ranges for a period of 15 months, where readings were made at a rate of 30 times per second, across approximately 240 sensors. We searched for intervals of 15 minutes where the sampling (specifically frequency, voltage phase angle and voltage magnitude) reflected an ambient conditions behavior and applied signal processing tools and statistical analyses on these periods.

By focusing on the frequency domain of the time series —obtained by its discrete-time Fourier transform— we observe anomalous peaks at 5 Hz, consistently present in a large portion of the buses across different days and moments of the day. We use Fourier filtering to isolate the contribution of these frequency modes or to suppress them.

The statistical tools used are principal component analysis (PCA) of the covariance (correlation) matrices, temporal auto-correlation and temporal cross-correlation of the normalized or filtered time series. From the first one, we conclude that the correlation matrix has low-rank, the largest 10 eigen-values account for more than 80% of the spectrum, with the first one representing more than 25%. We also notice that the eigen-values and eigen-vectors do not change dramatically over time, that is, they have a steady or a slow changing behavior during the analyzed quiet periods. Leading eigen-vectors are struc-

tured in such a way that sensors (associated with the components of the eigen-vectors) having the same value are geographically closed.

The correlation functions that we define account for the temporal correlation that might exist between a sensor's readings and its own past (auto-correlation), and also between a sensor's readings and the past of a different sensor (cross-correlation). The correlation functions that are obtained usually show a periodic behavior and, interestingly, it is mainly due to the modes around 5 Hz observed in the frequency domain of the series —when the modes near 5 Hz are suppressed, the correlation functions drop their value to zero.

The sensors with high auto-correlation are located in two specific areas of the map. This behavior is extended through the different observed periods. Moreover, the sensors located in these areas show high cross-correlation as well, which might indicated that they all function in a coordinated fashion.

As an ongoing work, we continue improving the filters that isolate different modes from the frequency domain of the times series. We will use the filtered times series to infer the branch admittance parameters across different transmission lines of the network.

Bibliography

[1] Report of WECC Joint Synchronized Information Subcommittee on Modes of Inter-Area Power Oscillations in Western Interconnection. https://www.wecc.biz/Reliability/WECCJSISModesofInter-AreaOscillations-2013-12-REV1.1.pdf, 2013. Accessed: 2018-09-11.

[2] Power System Oscillatory Behaviors: Sources, Characteristics, & Analyses. https://www.pnnl.gov/main/publications/external/technical_reports/PNNL-26375.pdf, 2017. Accessed: 2018-09-11.

[3] Reliability Guideline Forced Oscillation Monitoring & Mitigation. https://www.nerc.com/pa/RAPA/rg/ReliabilityGuidelines/Reliability_Guideline_-_Forced_Oscillations_-_2017.pdf, 2017. Accessed: 2018-09-11.

[4] Department of Energy: Big Data Analysis of Synchrophasor Data. https://grantbulletin.research.uiowa.edu/netl-big-data-analysis-synchrophasor-data, 2018.

[5] R. Ahuja, T. Magnanti, and J. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.

[6] E. D. Andersen and K. D. Andersen. *The Mosek Interior Point Optimizer for Linear Programming: An Implementation of the Homogeneous Algorithm*, pages 197–232. Springer US, Boston, MA, 2000.

[7] D. Apostolopoulou, P. W. Sauer, and A. D. Domínguez-García. Automatic Generation Control and Its Implementation in Real Time. In *2014 47th Hawaii International Conference on System Sciences*, pages 2444–2452, January 2014.

[8] I. Arel, D. Rose, and T. Karnowski. Deep Machine Learning - A New Frontier in Artificial Intelligence Research [Research Frontier]. *IEEE Computational Intelligence Mag.*, 5(4):13–18, November 2010.

[9] R. Arghandeh and Y. Zhou. *Big Data Application in Power Systems*. Elsevier, 2018.

[10] G. Bakke. *The Grid: The Fraying Wires between Americans and our Energy Future*. Bloomsbury USA, 2016.

[11] R. Baldick, K. A. Clements, Z. Pinjo-Dzigal, and P. W. Davis. Implementing Nonquadratic Objective Functions for State Estimation and Bad Data Rejection. *IEEE Transactions on Power Systems*, 12(1):376–382, February 1997.

[12] Y. Bengio. Learning Deep Architectures for AI. *Foundations and Trends in Machine Learning*, 2(1):1–127, January 2009.

[13] A. R. Bergen and V. Vittal. *Power Systems Analysis*. Prentice Hall, Upper Saddle River, NJ, Second edition, 2000.

[14] A. Bernstein, D. Bienstock, D. Hay, M. Uzunoglu, and G. Zussman. Power Grid Vulnerability to Geographically Correlated Failures Analysis and Control Implications. In *IEEE INFOCOM 2014 - IEEE Conference on Computer Communications*, pages 2634–2642, April 2014.

[15] S. Bhela, V. Kekatos, and S. Veeramachaneni. Enhancing Observability in Distribution Grids Using Smart Meter Data. *IEEE Transactions on Smart Grid*, 9(6):5953–5961, November 2018.

[16] S. Bhela, V. Kekatos, and S. Veeramachaneni. Smart inverter grid probing for learning loads: Part i - identifiability analysis. *IEEE Transactions on Power Systems*, to appear, 2019.

[17] S. Bhela, V. Kekatos, and S. Veeramachaneni. Smart inverter grid probing for learning loads: Part ii - probing injection design. *IEEE Transactions on Power Systems*, to appear, 2019.

[18] D. Bienstock. *Electrical Transmission System Cascades and Vulnerability: An Operations Research Viewpoint*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2015.

[19] D. Bienstock. Machine learning with PMU data. In *2017 NASPI Work Group Meeting, Gaithersburg, MD*, March 2017.

[20] D. Bienstock, M. Chertkov, and M. Escobar. Learning from ISO-scale PMU data stream, 2018.

[21] D. Bienstock, M. Chertkov, and M. Escobar. Learning from power system data stream: phasor-detective approach. *arXiv:1902.03223*, 2018.

[22] D. Bienstock, A. Shukla, and S. Yun. Non-Stationary Streaming PCA. *arXiv:1902.03223*, 2019.

[23] D. Bienstock and A. Verma. The $N - k$ Problem in Power Grids: New Models, Formulations, and Numerical Experiments. *SIAM Journal on Optimization*, 20(5):2352–2380, 2010.

[24] D. Bienstock and A. Verma. Strong NP-hardness of AC power flows feasibility. *arXiv:1512.07315*, December 2015.

[25] R. Bobba, K. Rogers, Q. Wang, H. Khurana, K. Nahrstedt, and T. Overbye. De-

tecting False Data Injection Attacks on DC State Estimation. In *Proceedings of the First Workshop on Secure Control Systems*, 2010.

[26] J. Carpentier. Contribution á l'étude du dispatching économique. *Bulletin de la Société Française des Électriciens*, 8(3):431–447, 1962.

[27] C. Chang and Z. Li. Recursive stochastic subspace identification for structural parameter estimation. In *Proceedings of SPIE*, volume 7292, March 2009.

[28] G. Chen, Z. Y. Dong, D. J. Hill, and Y. S. Xue. Exploring Reliable Strategies for Defending Power Systems Against Targeted Attacks. *IEEE Transactions on Power Systems*, 26(3):1000–1009, August 2011.

[29] L. Cheng, C. You, and L. Chen. Identification of Power Line Outages Based on PMU Measurements and Sparse Overcomplete Representation. In *2016 IEEE 17th International Conference on Information Reuse and Integration*, pages 343–349, July 2016.

[30] T. V. Cutsem, M. Ribbens-Pavella, and L. Mili. Hypothesis Testing Identification: A New Method For Bad Data Analysis In Power System State Estimation. *IEEE Transactions on Power Apparatus and Systems*, PAS-103(11):3239–3252, November 1984.

[31] N. Dahal, R. L. King, and V. Madani. Online dimension reduction of synchrophasor data. In *PES T D 2012*, pages 1–7, May 2012.

[32] G. Dán and H. Sandberg. Stealth Attacks and Protection Schemes for State Estimators in Power Systems. In *2010 First IEEE International Conference on Smart Grid Communications*, pages 214–219, October 2010.

[33] R. Deng, P. Zhuang, and H. Liang. CCPA: Coordinated Cyber-Physical Attacks and Countermeasures in Smart Grid. *IEEE Transactions on Smart Grid*, 8(5):2420–

2430, September 2017.

[34] Electricity-Information Sharing and Analysis Center (E-ISAC). Analysis of the Cyber Attack on the Ukrainian Power Grid. Technical report, SANS Industrial Control Systems, March 2016.

[35] M. Esmalifalak, H. Nguyen, and R. Zheng. Stealth False Data Injection using Independent Component Analysis in Smart Grid. In *2011 IEEE International Conference on Smart Grid Communications*, pages 244–248, October 2011.

[36] G. Frigo, C. Narduzzi, D. Colangelo, M. Pignati, and M. Paolone. Definition and assessment of reference values for PMU calibration in static and transient conditions. In *2016 IEEE International Workshop on Applied Measurements for Power Systems (AMPS)*, pages 1–6, September 2016.

[37] P. Gao, M. Wang, J. H. Chow, S. G. Ghiocel, B. Fardanesh, G. Stefopoulos, and M. P. Razanousky. Identification of Successive "Unobservable" Cyber Data Attacks in Power Systems Through Matrix Decomposition. *IEEE Transactions on Signal Processing*, 64(21):5557–5570, November 2016.

[38] P. Gao, M. Wang, S. G. Ghiocel, J. H. Chow, B. Fardanesh, and G. Stefopoulos. Missing Data Recovery by Exploiting Low-Dimensionality in Power System Synchrophasor Measurements. *IEEE Transactions on Power Systems*, 31(2):1006–1013, March 2016.

[39] M. Garcia, T. Catanach, S. V. Wiel, R. Bent, and E. Lawrence. Line Outage Localization Using Phasor Measurement Data in Transient State. *IEEE Transactions on Power Systems*, 31(4):3019–3027, July 2016.

[40] A. Giani, E. Bitar, M. Garcia, M. McQueen, P. Khargonekar, and K. Poolla. Smart Grid Data Integrity Attacks. *IEEE Transactions on Smart Grid*, 4(3):1244–1253,

September 2013.

[41] J. D. Glover, M. S. Sarma, and T. J. Overbye. *Power Systems Analysis and Design*. Cengage Learning, Stamford, CT, Fourth edition, 2008.

[42] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. The MIT Press, 2016.

[43] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2018. http://www.gurobi.com.

[44] F. J. Harris. On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform. *Proceedings of the IEEE*, 66(1):51–83, January 1978.

[45] J. F. Hauer, C. J. Demeure, and L. L. Scharf. Initial results in Prony analysis of power system response signals. *IEEE Transactions on Power Systems*, 5(1):80–89, February 1990.

[46] Y. Huang, S. Werner, J. Huang, N. Kashyap, and V. Gupta. State Estimation in Electric Power Grids: Meeting New Challenges Presented by the Requirements of the Future Grid. *IEEE Signal Processing Magazine*, 29(5):33–43, September 2012.

[47] G. Hug and J. A. Giampapa. Vulnerability Assessment of AC State Estimation With Respect to False Data Injection Cyber-Attacks. *IEEE Transactions on Smart Grid*, 3(3):1362–1370, September 2012.

[48] International Bureau of Weights and Measures (BIPM). The International System of Units (SI). Brochure, 2006.

[49] M. Kezunovic, L. Xie, and S. Grijalva. The role of big data in improving power system operation and protection. In *2013 IREP Symposium Bulk Power System Dynamics and Control - IX Optimization, Security and Control of the Emerging Power Grid*, pages 1–9, August 2013.

[50] J. Kim and L. Tong. On Topology Attack of a Smart Grid: Undetectable Attacks and Countermeasures. *IEEE Journal on Selected Areas in Communications*, 31(7):1294–1305, July 2013.

[51] J. Kim, L. Tong, and R. J. Thomas. Subspace Methods for Data Attack on State Estimation: A Data Driven Approach. *IEEE Transactions on Signal Processing*, 63(5):1102–1114, March 2015.

[52] T. T. Kim and H. V. Poor. Strategic Protection Against Data Injection Attacks on Power Grids. *IEEE Transactions on Smart Grid*, 2(2):326–333, June 2011.

[53] B. J. Kirby, J. Dyer, C. Martinez, R. A. Shoureshi, R. Guttromson, and J. Dale. Frequency Control Concerns in the North American Electric Power System. Technical Report ORNL/TM-2003/41, U.S. Department of Energy, December 2002.

[54] B. Kocuk, S. S. Dey, and X. A. Sun. Inexactness of SDP Relaxation and Valid Inequalities for Optimal Power Flow. *IEEE Transactions on Power Systems*, 31(1):642–651, January 2016.

[55] B. Kocuk, S. S. Dey, and X. A. Sun. Strong SOCP Relaxations for the Optimal Power Flow Problem. *Operations Research*, 64(6):1177–1196, 2016.

[56] Y. LeCun, Y. Bengio, and G. Hinton. Deep Learning. *Nature*, 521(7553):436–444, May 2015.

[57] W. Li, M. Wang, and J. Chow. Real-Time Event Identification Through Low-Dimensional Subspace Characterization of High-Dimensional Synchrophasor Data. *IEEE Trans. on Power Systems*, 33(5):4937–4947, September 2018.

[58] Z. Li, M. Shahidehpour, A. Alabdulwahab, and A. Abusorrah. Bilevel Model for Analyzing Coordinated Cyber-Physical Attacks on Power Systems. *IEEE Transactions on Smart Grid*, 7(5):2260–2272, September 2016.

[59] J. Liang, O. Kosut, and L. Sankar. Cyber Attacks on AC State Estimation: Unobservability and Physical Consequences. In *2014 IEEE PES General Meeting — Conference Exposition*, pages 1–5, July 2014.

[60] J. Liang, L. Sankar, and O. Kosut. Vulnerability Analysis and Consequences of False Data Injection Attack on Power System State Estimation. *IEEE Transactions on Power Systems*, 31(5):3864–3872, September 2016.

[61] J. Lin and H. Pan. A Static State Estimation Approach Including Bad Data Detection and Identification in Power Systems. In *2007 IEEE Power Engineering Society General Meeting*, pages 1–7, June 2007.

[62] G. Liu, J. Quintero, and V. M. Venkatasubramanian. Oscillation monitoring system based on wide area synchrophasors in power systems. In *2007 iREP Symposium - Bulk Power System Dynamics and Control - VII. Revitalizing Operational Reliability*, pages 1–13, August 2007.

[63] X. Liu and Z. Li. Local Load Redistribution Attacks in Power Systems With Incomplete Network Information. *IEEE Transactions on Smart Grid*, 5(4):1665–1676, July 2014.

[64] Y. Liu, P. Ning, and M. K. Reiter. False Data Injection Attacks Against State Estimation in Electric Power Grids. In *Proceedings of the 16th ACM Conference on Computer and Communications Security*, CCS '09, pages 21–32, New York, NY, USA, 2009. ACM.

[65] A. Lokhov, M. Vuffray, D. Shemetov, D. Deka, and M. Chertkov. Online Learning of Power Transmission Dynamics. *PSCC 2018, arXiv:1710.10021*, 2017.

[66] N. M. Manousakis, G. N. Korres, and P. S. Georgilakis. Taxonomy of PMU Placement Methodologies. *IEEE Transactions on Power Systems*, 27(2):1070–1077, May

2012.

[67] D. K. Molzahn and J. Wang. Detection and Characterization of Intrusions to Network Parameter Data in Electric Power Systems. *IEEE Transactions on Smart Grid*, to appear, 2018.

[68] A. Monticelli and A. Garcia. Reliable Bad Data Processing for Real-Time State Estimation. *IEEE Transactions on Power Apparatus and Systems*, PAS-102(5):1126–1139, May 1983.

[69] A. Monticelli, F. F. Wu, and M. Yen. Mutiple Bad Data Identification for State Estimation by Combinatorial Optimization. *IEEE Transactions on Power Delivery*, 1(3):361–369, July 1986.

[70] K. Narendra, D. Rangana, and A. Rajapakse. Dynamic Performance Evaluation and Testing of Phasor Measurement Unit (PMU) as per IEEE C37.118.1 Standard. In *Doble Client Committee Meetings & International Protection Testing Users Group (PTUG)*, October 2010.

[71] A. V. Oppenheim, R. W. Schafer, and J. R. Buck. *Discrete-Time Signal Processing*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, Second edition, 1999.

[72] A. Pinar, J. Meza, V. Donde, and B. Lesieutre. Optimization Strategies for the Vulnerability Analysis of the Electric Power Grid. *SIAM Journal on Optimization*, 20(4):1786–1810, February 2010.

[73] M. A. Rahman and H. Mohsenian-Rad. False Data Injection Attacks with Incomplete Information Against Smart Power Grids. In *2012 IEEE Global Communications Conference*, pages 3153–3158, December 2012.

[74] S. Roy and B. Lesieutre. Frequency Band Decomposition of a Dynamic Persistence Measure Using Ambient Synchrophasor Data. In *2018 Power Systems Computation*

Conference (PSCC), pages 1–7, June 2018.

[75] S. Sarmadi and V. Venkatasubramanian. Inter-Area Resonance in Power Systems From Forced Oscillations. *IEEE Trans. on Power Systems*, 31(1):378–386, January 2016.

[76] J. Schmidhuber. Deep Learning in Neural Networks. *Neural Networks*, 61(C):85–117, January 2015.

[77] F. C. Schweppe. Power System Static-State Estimation, Part III: Implementation. *IEEE Transactions on Power Apparatus and Systems*, PAS-89(1):130–135, January 1970.

[78] F. C. Schweppe and D. B. Rom. Power System Static-State Estimation, Part II: Approximate Model. *IEEE Transactions on Power Apparatus and Systems*, PAS-89(1):125–130, January 1970.

[79] F. C. Schweppe and J. Wildes. Power System Static-State Estimation, Part I: Exact Model. *IEEE Transactions on Power Apparatus and Systems*, PAS-89(1):120–125, January 1970.

[80] S. Soltan, M. Yannakakis, and G. Zussman. Power Grid State Estimation Following a Joint Cyber and Physical Attack. *IEEE Transactions on Control of Network Systems*, 5(1):499–512, March 2018.

[81] S. Soltan, M. Yannakakis, and G. Zussman. REACT to Cyber Attacks on Power Grids. *IEEE Transactions on Network Science and Engineering*, to appear, 2018.

[82] S. Soltan and G. Zussman. Power Grid State Estimation after a Cyber-Physical Attack under the AC Power Flow Model. In *2017 IEEE Power Energy Society General Meeting*, pages 1–5, July 2017.

[83] S. Soltan and G. Zussman. EXPOSE the Line Failures following a Cyber-Physical Attack on the Power Grid. *IEEE Transactions on Control of Network Systems*, to appear, 2018.

[84] Y. Tang and G. N. Stenbakken. Calibration of phasor measurement unit at NIST. In *2012 Conference on Precision electromagnetic Measurements*, pages 414–415, July 2012.

[85] J. E. Tate and T. J. Overbye. Line Outage Detection Using Phasor Angle Measurements. *IEEE Transactions on Power Systems*, 23(4):1644–1652, November 2008.

[86] J. E. Tate and T. J. Overbye. Double Line Outage Detection Using Phasor Angle Measurements. In *2009 IEEE Power Energy Society General Meeting*, pages 1–5, July 2009.

[87] A. Teixeira, S. Amin, H. Sandberg, K. H. Johansson, and S. S. Sastry. Cyber Security Analysis of State Estimators in Electric Power Systems. In *49th IEEE Conference on Decision and Control (CDC)*, pages 5991–5998, December 2010.

[88] D. Trudnowski. Properties of the Dominant Inter-Area Modes in the WECC Interconnect. https://www.wecc.biz/Reliability/WECCmodesPaper130113Trudnowski.pdf, 2012. Accessed: 2018-09-11.

[89] D. Trudnowski and J. Pierre. *Signal Processing Methods for Estimating Small-Signal Dynamic Properties from Measured Responses*, pages 1–36. Springer US, Boston, MA, 2009.

[90] D. Van Hertem, J. Verboomen, K. Purchala, R. Belmans, and W. L. Kling. Usefulness of DC Power Flow for Active Power Flow Analysis with Flow Controlling Devices. In *The 8th IEE International Conference on AC and DC Power Transmission*, pages 58–62, March 2006.

[91] O. Vuković, K. C. Sou, G. Dán, and H. Sandberg. Network-layer Protection Schemes against Stealth Attacks on State Estimators in Power Systems. In *2011 IEEE International Conference on Smart Grid Communications*, pages 184–189, October 2011.

[92] A. Wächter and L. T. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106(1):25–57, March 2006.

[93] L. Xie, Y. Chen, and P. R. Kumar. Dimensionality Reduction of Synchrophasor Data for Early Event Detection: Linearized Analysis. *IEEE Transactions on Power Systems*, 29(6):2784–2794, November 2014.

[94] L. Xie, Y. Chen, and P. R. Kumar. Dimensionality Reduction of Synchrophasor Data for Early Event Detection: Linearized Analysis. *IEEE Transactions on Power Systems*, 29(6):2784–2794, November 2014.

[95] Y. Yuan, Z. Li, and K. Ren. Modeling Load Redistribution Attacks in Power Systems. *IEEE Transactions on Smart Grid*, 2(2):382–390, June 2011.

[96] J. Zhang and L. Sankar. Physical System Consequences of Unobservable State-and-Topology Cyber-Physical Attacks. *IEEE Transactions on Smart Grid*, 7(4):2016–2025, July 2016.

[97] Y. Zhao, J. Chen, A. Goldsmith, and H. V. Poor. Identification of Outages in Power Systems With Uncertain States and Optimal Sensor Locations. *IEEE Journal of Selected Topics in Signal Processing*, 8(6):1140–1153, December 2014.

[98] Y. Zhao, A. Goldsmith, and H. V. Poor. On PMU Location Selection for Line Outage Detection in Wide-area Transmission Networks. In *2012 IEEE Power and Energy Society General Meeting*, pages 1–8, July 2012.

[99] N. Zhou, D. J. Trudnowski, J. W. Pierre, and W. A. Mittelstadt. Electromechanical Mode Online Estimation Using Regularized Robust RLS Methods. *IEEE Transactions on Power Systems*, 23(4):1670–1680, November 2008.

[100] H. Zhu and G. B. Giannakis. Sparse Overcomplete Representations for Efficient Identification of Power Line Outages. *IEEE Transactions on Power Systems*, 27(4):2215–2224, November 2012.

[101] R. D. Zimmerman, C. E. Murillo-Sanchez, and R. J. Thomas. MATPOWER: Steady-State Operations, Planning, and Analysis Tools for Power Systems Research and Education. *IEEE Transactions on Power Systems*, 26(1):12–19, February 2011.

Appendices

Appendix

Correlation Plots

A.1 Auto-Correlation: Voltage Phase Angle

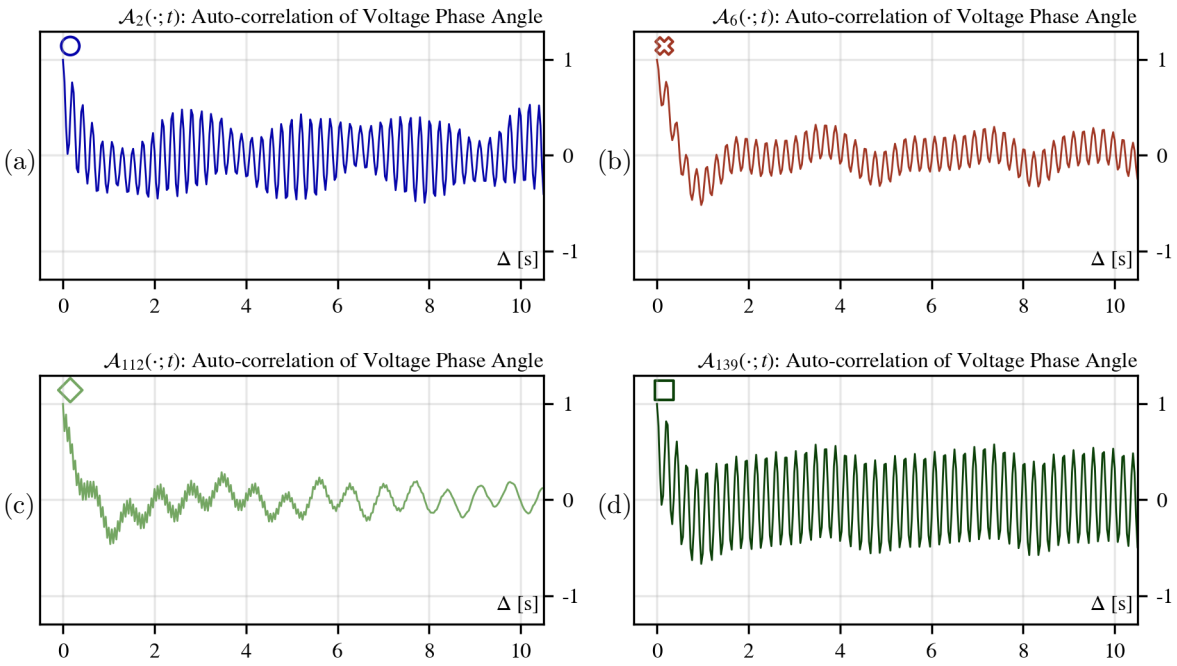

Figure A.1: Auto-correlation functions for voltage phase angle for PMU's (a) $k = 2$, (b) $k = 6$, (c) $k = 112$, and (d) $k = 139$; at t=21:29:00 on July 30, 2013. The correlation matrices are constructed with the normalized time series $\hat{m}^{(s)}(\cdot; S)$, $S = 30$, $T = 5400$.

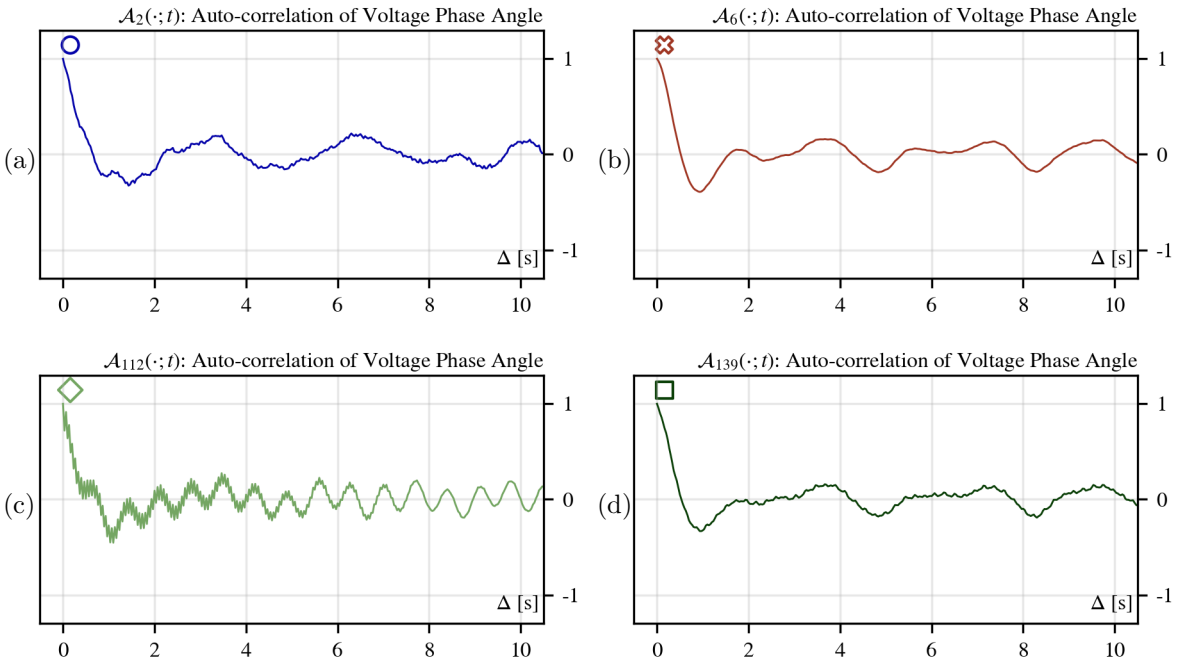

Figure A.2: Auto-correlation functions for voltage phase angle for PMU's (a) $k = 2$, (b) $k = 6$, (c) $k = 112$, and (d) $k = 139$; at t=21:29:00 on July 30, 2013. The correlation matrices are constructed with the normalized time series $\hat{f}^{(s)}[\tilde{\varphi}^{\mathrm{BSF5}}](\cdot;S)$ that has been filtered by a band-stop Fourier filter, $S = 30$, $T = 5400$.

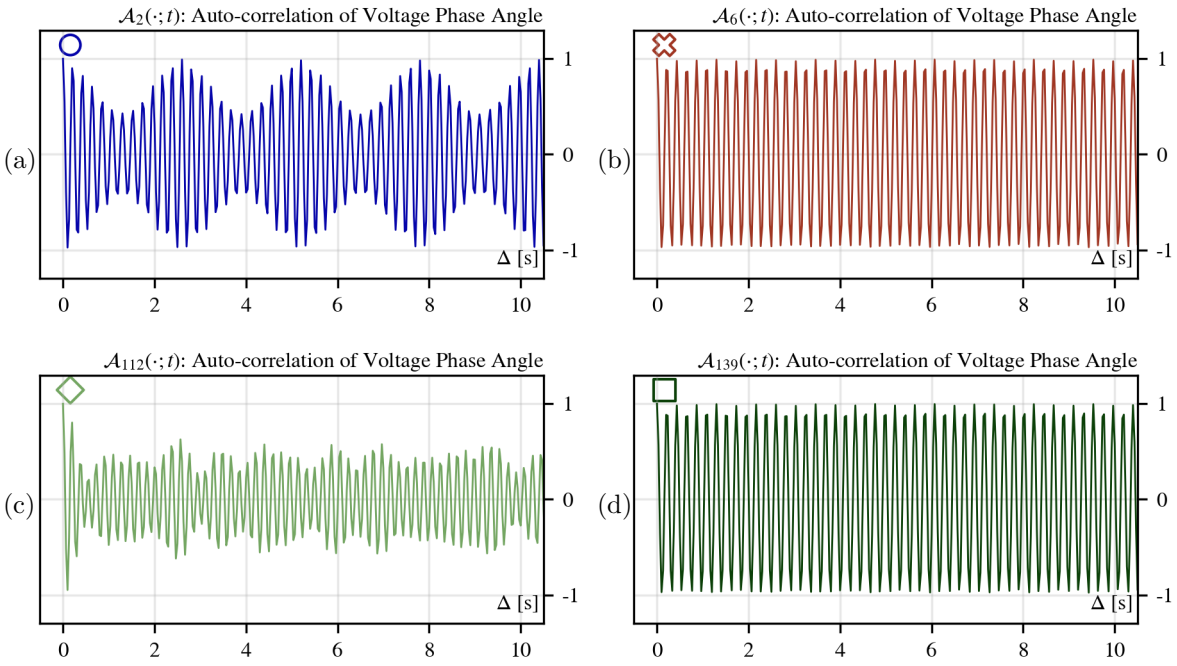

Figure A.3: Auto-correlation functions for voltage phase angle for PMU's (a) $k = 2$, (b) $k = 6$, (c) $k = 112$, and (d) $k = 139$; at t=21:29:00 on July 30, 2013. The correlation matrices are constructed with the normalized time series $\hat{f}^{(s)}[\tilde{\varphi}^{\mathrm{BPF5}}](\cdot;S)$ that has been filtered by a band-pass Fourier filter, $S = 30$, $T = 5400$.

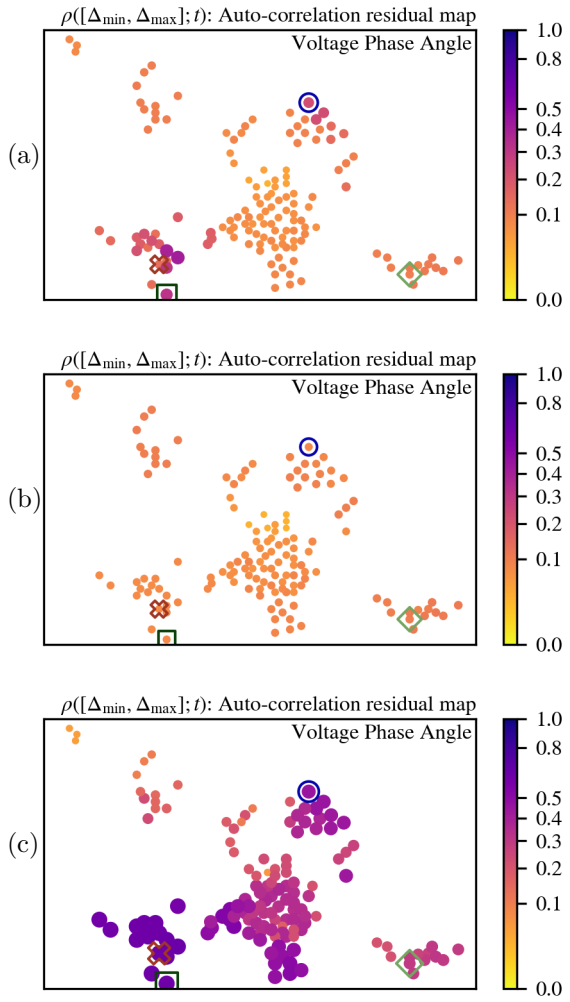

Figure A.4: Residue of auto-correlation functions for voltage phase angle using (a) $\hat{m}^{(s)}(\cdot; S)$, (b) $\hat{f}^{(s)}[\tilde{\varphi}^{\mathrm{BSF5}}](\cdot; S)$, and (c) $\hat{f}^{(s)}[\tilde{\varphi}^{\mathrm{BPF5}}](\cdot; S)$; at t=21:29:00 on July 30, 2013, with $S = 30$, $T = 5400$, $\Delta_{\min} = 30$, and $\Delta_{\max} = 1800$. Geometrical figures show the position of the sensors depicted in Figures A.1–A.3.

A.2 Auto-Correlation: Voltage Magnitude

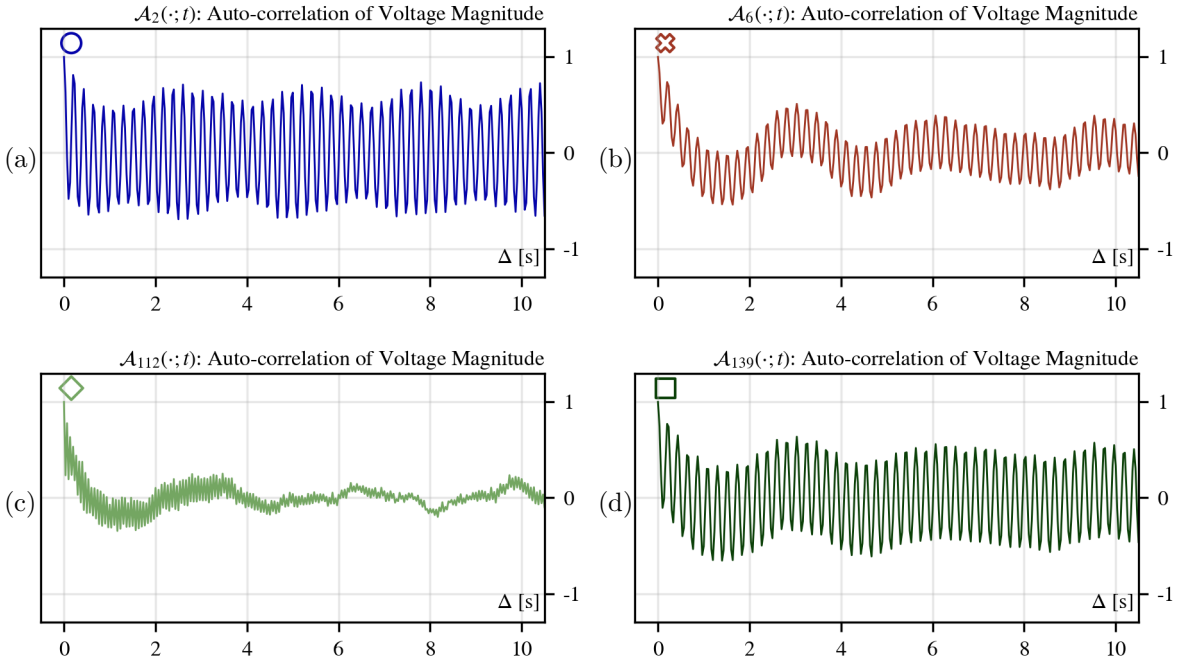

Figure A.5: Auto-correlation functions for voltage magnitude for PMU's (a) $k = 2$, (b) $k = 6$, (c) $k = 112$, and (d) $k = 139$; at t=21:29:00 on July 30, 2013. The correlation matrices are constructed with the normalized time series $\hat{m}^{(s)}(\cdot; S)$, $S = 30$, $T = 5400$.

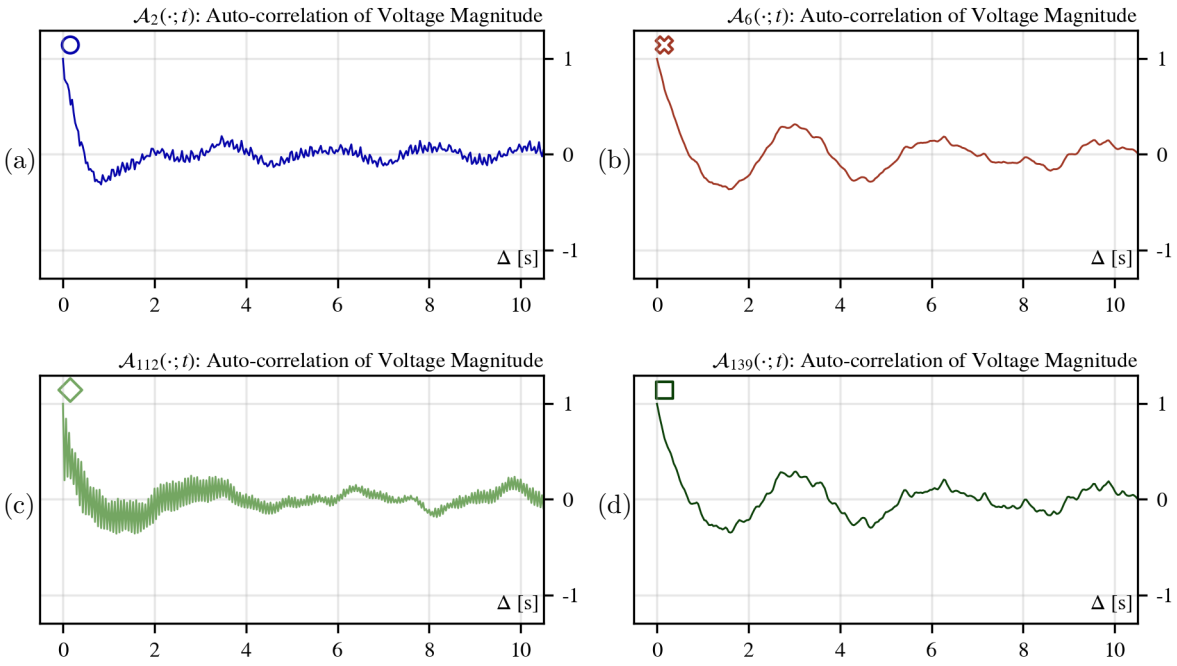

Figure A.6: Auto-correlation functions for voltage magnitude for PMU's (a) $k = 2$, (b) $k = 6$, (c) $k = 112$, and (d) $k = 139$; at t=21:29:00 on July 30, 2013. The correlation matrices are constructed with the normalized time series $\hat{f}^{(s)}[\tilde{\varphi}^{\mathrm{BSF5}}](\cdot;S)$ that has been filtered by a band-stop Fourier filter, $S = 30$, $T = 5400$.

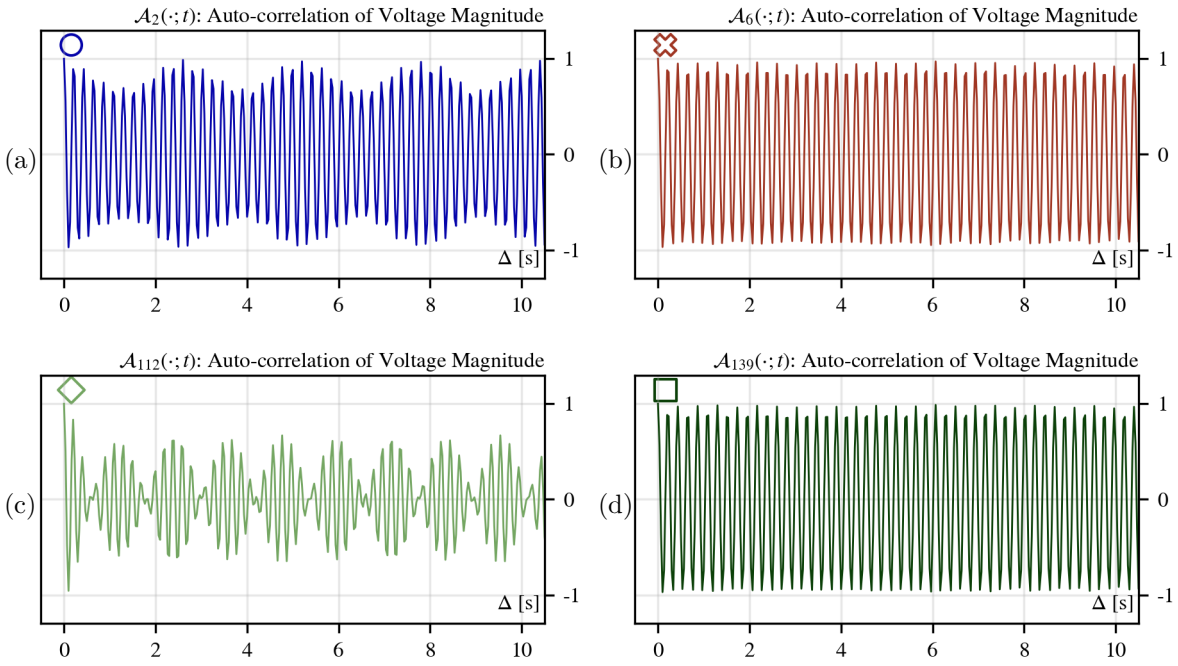

Figure A.7: Auto-correlation functions for voltage magnitude for PMU's (a) $k = 2$, (b) $k = 6$, (c) $k = 112$, and (d) $k = 139$; at t=21:29:00 on July 30, 2013. The correlation matrices are constructed with the normalized time series $\hat{f}^{(s)}[\tilde{\varphi}^{\mathrm{BPF5}}](\cdot;S)$ that has been filtered by a band-pass Fourier filter, $S = 30$, $T = 5400$.

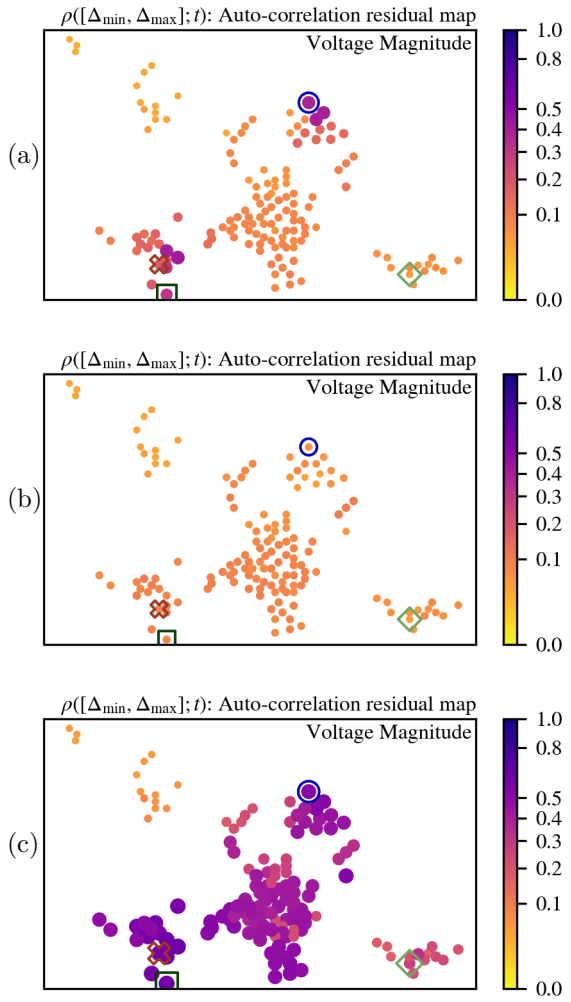

Figure A.8: Residue of auto-correlation functions for voltage magnitude using (a) $\hat{m}^{(s)}(\cdot; S)$, (b) $\hat{f}^{(s)}[\tilde{\varphi}^{\mathrm{BSF5}}](\cdot; S)$, and (c) $\hat{f}^{(s)}[\tilde{\varphi}^{\mathrm{BPF5}}](\cdot; S)$; at t=21:29:00 on July 30, 2013, with $S = 30$, $T = 5400$, $\Delta_{\min} = 30$, and $\Delta_{\max} = 1800$. Geometrical figures show the position of the sensors depicted in Figures A.5–A.7.

A.3 Cross-Correlation: Voltage Phase Angle

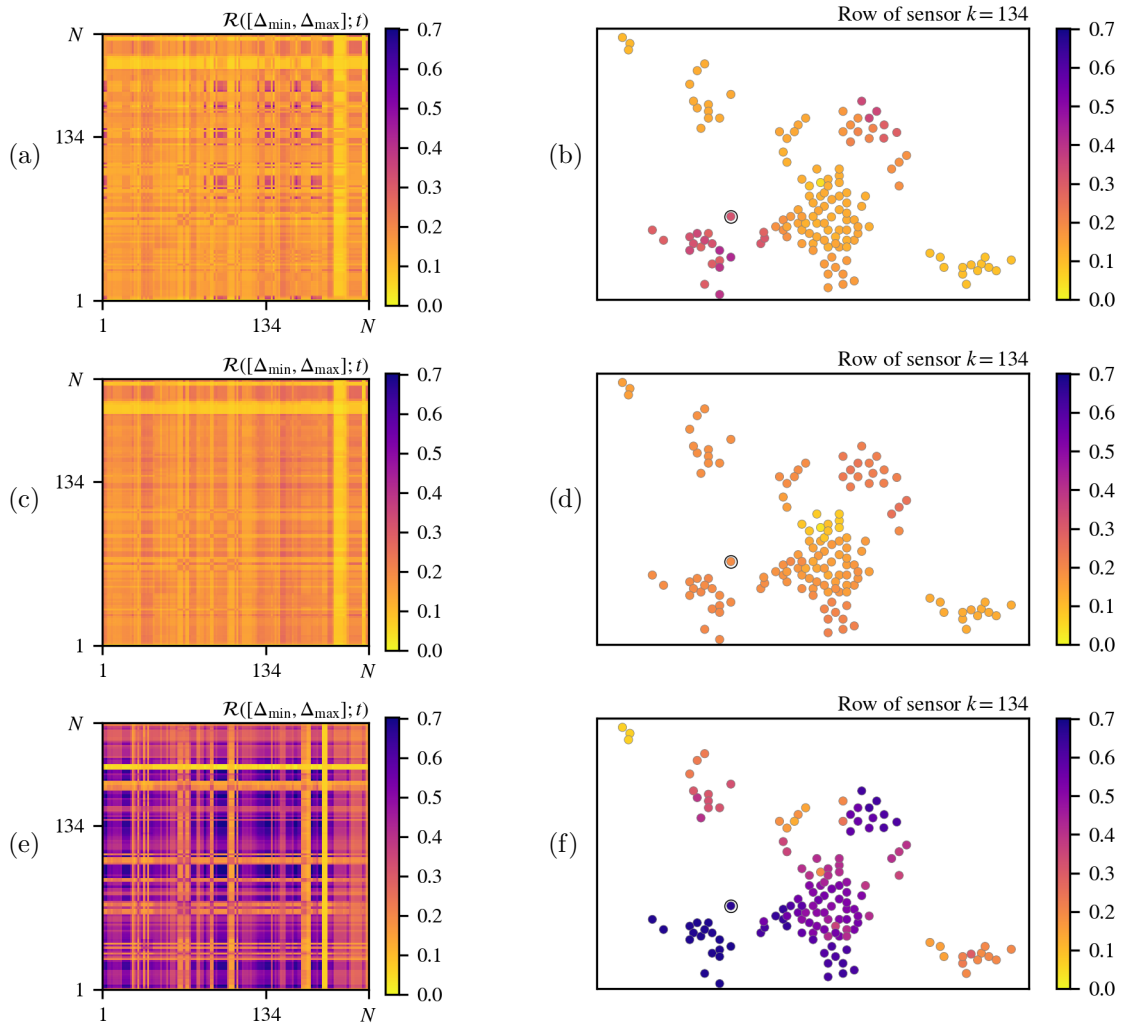

Figure A.9: Cross-correlation residual matrix (on the left plots) and cross-correlation residue for voltage phase angle, geographically located, between sensor $k = 134$ —indicated with an extra circle around it— and the remaining sensors (on the right plots) using $\hat{m}^{(s)}(\cdot; S)$ in (a) and (b), $\hat{f}^{(s)}[\tilde{\varphi}^{\text{BSF5}}](\cdot; S)$ in (c) and (d), and $\hat{f}^{(s)}[\tilde{\varphi}^{\text{BPF5}}](\cdot; S)$ in (e) and (f); at t=21:29:00 on July 30, 2013, with $S = 30$, $T = 5400$, $\Delta_{\min} = 30$, and $\Delta_{\max} = 1800$.

A.4 Cross-Correlation: Voltage Magnitude

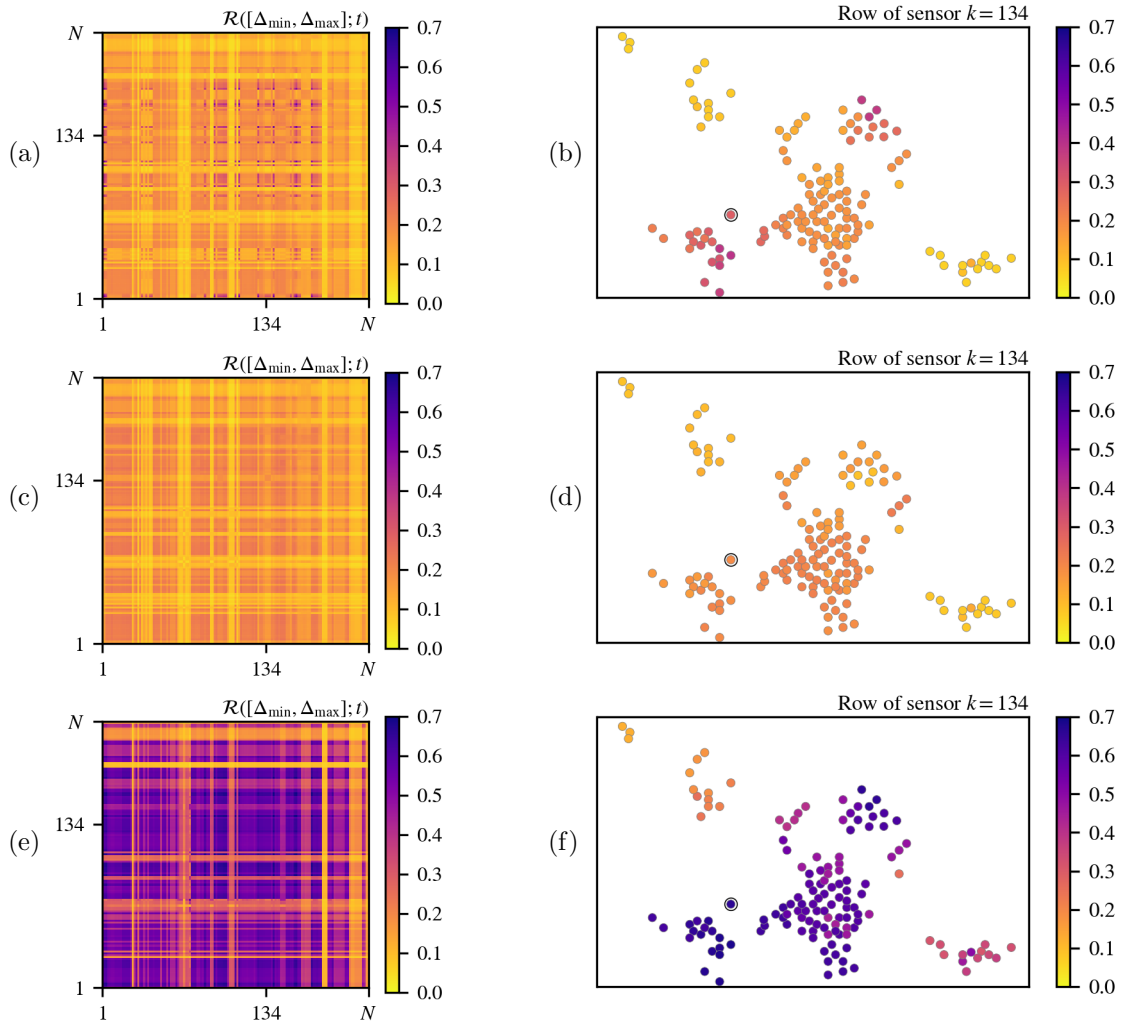

Figure A.10: Cross-correlation residual matrix (on the left plots) and cross-correlation residue for voltage magnitude, geographically located, between sensor $k = 134$ —indicated with an extra circle around it— and the remaining sensors (on the right plots) using $\hat{m}^{(s)}(\cdot; S)$ in (a) and (b), $\hat{f}^{(s)}[\tilde{\varphi}^{\mathrm{BSF5}}](\cdot; S)$ in (c) and (d), and $\hat{f}^{(s)}[\tilde{\varphi}^{\mathrm{BPF5}}](\cdot; S)$ in (e) and (f); at t=21:29:00 on July 30, 2013, with $S = 30$, $T = 5400$, $\Delta_{\min} = 30$, and $\Delta_{\max} = 1800$.