Phasor-Based Control on Islanded Networks

Kyle Brady

Electrical Engineering and Computer Sciences
University of California, Berkeley

Technical Report No. UCB/EECS-2021-233
http://www2.eecs.berkeley.edu/Pubs/TechRpts/2021/EECS-2021-233.html

December 1, 2021

Phasor-Based Control on Islanded Networks

by

Kyle Brady

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering - Electrical Engineering and Computer Sciences

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Adjunct Professor Alexandra von Meier, Co-chair
Professor Seth Sanders, Co-chair
Associate Professor Scott Moura

Fall 2020

Phasor-Based Control on Islanded Networks

Abstract

Phasor-Based Control on Islanded Networks

by

Kyle Brady

Doctor of Philosophy in Engineering - Electrical Engineering and Computer Sciences

University of California, Berkeley

Adjunct Professor Alexandra von Meier, Co-chair

Professor Seth Sanders, Co-chair

The advancement of sensing technologies for the power grid has allowed the development of new strategies for the control of distributed energy resources (DERs). In particular, the emergence of phasor measurement units (PMUs) designed for deployment at the distribution level has presented an exciting opportunity. These PMUs have enabled the development of phasor-based control (PBC), a strategy that formulates DER power dispatch in terms of voltage phasor targets to be tracked by local controllers.

This dissertation focuses on the optimal power flow (OPF) component of PBC's supervisory control layer, which has previously been conceptualized and demonstrated in simulation on distribution networks. We expand its applicability to medium-voltage minigrids and microgrids operating in island mode, networks where PBC has the potential to deliver important benefits.

The work is carried out in two stages. After a discussion of PBC and other relevant background topics, we address one of the primary challenges to PBC at the medium voltage level: the need for extreme accuracy in the supervisory controller's generation of phasor targets. This accuracy is achieved through an adaptation of an iterative OPF methodology that refines a linearized model of power flow through successive exchanges with a nonlinear solver. We discuss the changes that were made to both linear model and nonlinear solver, as well as the determination of phasor targets on networks that include tap-changing transformers and other realistic equipment. The accuracy of the adapted iterative method is then shown in simulation.

The second stage of the work covers the extension of our OPF implementation to islanded systems. We present a strategy for the treatment of the slack bus used by our nonlinear solver and apply it to several test cases in simulation. We then analyze a specific case in

which our iterative solution method fails, and demonstrate the use of a penalty factor in our linearized OPF formulation as a means of overcoming that failure. A full, end-to-end implementation of PBC's supervisory layer is then proposed and tested on a number of DER distributions at different feeder penetration levels.

We end with a presentation of data relevant to instrument-transformer-induced error in PMU measurements. This final portion stands alone from the primary work of the dissertation, but remains highly relevant to PBC. From an experimental deployment of two PMUs measuring an identical distribution-grid voltage, we determine a ratio of the errors induced by their individual potential transformers. Monitoring this quantity over the course of a year allows us to track the drift in those induced errors over time, motivating a discussion of the expected impact of error drift on PBC and the frequency with which transformers will need to be recalibrated in operational settings.

To Liz, who has earned it many times over.

Contents

List of Figures

List of Tables

Acknowledgments

Throughout my years at Berkeley, many people have offered me their time, expertise, and friendship. I am in their debt.

First and foremost, I thank my advisor, Dr. Alexandra "Sascha" von Meier. An incredible mentor, Dr. von Meier encouraged me to explore my areas of interest, always providing timely guidance to help me structure those interests into focused, viable projects.

Ever approachable and engaged, Dr. von Meier exemplifies the art of wearing great learning lightly. She can draw on an incredible depth of knowledge to address complicated issues, often in real time and while under pressure from multiple simultaneous projects. She does it with such effortless grace and modesty, though, that she puts everyone from first-year students to program directors at ease. In engineering, rock climbing, and a considered approach to learning, I could not have asked for a better model.

I am also deeply appreciative of the professors who have provided additional guidance during my time in the program. In particular, I am grateful to Drs. Seth Sanders, Scott Moura, and Claire Tomlin for the time and energy they have expended serving on both my dissertation and qualifying exam committees. The insights they have offered have been invaluable in the writing of this work, and their readiness to fit me into extremely busy schedules (particularly in a year of unprecedented logistical obstacles) has been beyond generous.

I thank the students, instructors, and collaborators who have made my experience here both rewarding and genuinely fun. As a non-exhaustive list, I want to mention Ahmed Badruzzaman, Alex McEachern, Alexander Fiore, Aminy Ostfeld, Antoine Lesage-Landry, Bernhard Boser, Brian Johnson, Carlo Muscas, Christine Rosen, Chris Yarp, Ciaran Roberts, Dan Arnold, Daniel Gerber, Dexter Scobee, Duncan Callaway, Emma Stewart, Gabe Fierro, Gabriel Colon-Reyes, Hari Prasanna Das, Ioannis Konstantakopoulous, Jaimie Swartz, Jalel Sager, Jason Poon, Jasper Pakshong, Jonathan Lee, Jorge Elizondo, Jose Lara, Josep Guerrero, Justin Yim, Kasra Nowrouzi, Keith Moffat, Kelly Fernandez, Laura Mehrmanesh, Laurel Dunn, Liz Ratnam, Lucas Spangher, Maruf Ahmed, Michael Andersen, Michael Sankur, Miles Rusch, Ming Jin, Mohini Bariya, Moustafa Abdelbaky, Omid Ardakanian, Paolo Castello, Paolo Pegoraro, Paroma Chatterjee, Phillippe Phanivong, Reza Arghandeh, Roberto Puddu, Roel Dobbe, Sidney Buchbinder, and Vishnu Murthy.

Another huge thank you goes out to the EECS department staff, with a special mention for Shirley Salanio. Shirley's heroic efforts on behalf of students have continually amazed me; it is a rare problem that cannot be solved by asking her advice. On numerous occasions, Shirley has gone out of her way to make life easier for me and my colleagues, and she does it all with such warmth and cheerfulness that you cannot help feeling happier every time you pass through her office. The graduate program and Berkeley as a whole are truly lucky to have her.

Outside of Berkeley, I want to thank the team at Zola Electric. The opportunity to work with them was a chance to experience firsthand the challenges facing those building solutions to difficult problems. My supervisor, Claudio Vergara, deserves extra thanks for all his help

and support throughout my time at the company. I spent the entirety of the year we worked together in awe of his knowledge, drive, and intellectual curiosity. This dissertation owes another huge debt to his assistance in structuring my thoughts.

Lastly, I offer gratitude to my family. My wife, Liz, has been more help throughout these years than I can possibly express. She's been a companion, a sounding board, a mostly willing practice presentation audience, and a good sport during my debugging sessions. She has put up with everything that comes with having a grad student in the house, and she's done it while earning her MBA and launching herself into a career in an incredibly challenging field. Every day I am more convinced that asking her to marry me is the best decision I will ever make. My parents, too, have been cheering me on at every step. They have helped me in a million ways, large and small, and it is because of their unconditional support that I am here now. My brother and his family I credit with keeping me grounded, and my grandmother is forever my model of "never give up, never give in." All my aunts, uncles, cousins, adopted family, and family of friends have been unfailingly supportive through everything I have done, and as I move into this next chapter of life I do it with total confidence and gratitude for the people who will be coming with me.

Nomenclature

3PH4W	3-Phase, 4-Wire
AC	Alternating Current
BFS	Backward-Forward Sweep
CERTS	Consortium for Electric Reliability Technology Solutions
CHP	Combined Heat and Power
CIM	Current Injection Method
CT	Current Transformer
DC	Direct Current
DER	Distributed Energy Resource
EV	Electric Vehicle
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
IEEE13	IEEE 13-Node Test Feeder
ISO	Independent System Operator
L-PBC	Local Phasor-Based Controller
LMP	Locational Marginal Price/Pricing
LUPFM	Linear Unbalanced Power Flow Model
LTC	Load Tap Changer
NR	Newton-Raphson
OPF	Optimal Power Flow

OLS	Ordinary Least Squares
PBC	Phasor-Based Control
PMU	Phasor Measurement Unit
PT	Potential Transformer
PV	Photovoltaic
RCAC	Retrospective Cost-Adaptive Control
RCF	Ratio Correction Factor
RMS	Root Mean Square
S-PBC	Supervisory Phasor-Based Controller
SDG	(United Nations) Sustainable Development Goal
SDP	Semi-Definite Program/Programming
SHS	Solar Home Sytem
SOCP	Second-Order Cone Program/Programming
SPIDERS	Smart Power Infrastructure Demonstration for Energy Reliability and Security
VA,VAR	Volt-Amperes, Volt-Ampere Reactive
WAMS	Wide Area Monitoring System
WLS	Weighted Least Squares

Chapter 1

Introduction

This introduction, along with Chapter 2, will lay the groundwork for the dissertation. The introduction's primary focus will be motivating the research topic: the extension of a novel, phasor-based control (PBC) strategy to allow for the coordination of single-phase generation sources on islanded, medium-voltage electrical networks.

To begin, this chapter briefly describes the structure of the electric grid, specifies the level of operation within that structure at which PBC is relevant, and introduces necessary background concepts and terminology. It ends with a high-level description of the work to follow and a guide to the organization of the dissertation.

1.1 From Centralized to Distributed Power Generation

The power grid as a concept has been with us since the end of the 19th century. From its roots in local DC power generation plants in the United States, electrical grids have expanded into continent-spanning networks of interconnected generators serving the majority of the world's population. As might be expected from a technical innovation that has existed at such a massive scale for over a hundred years, power grids are not homogenous. In different regions, implementation of "the grid" reflects a unique history of technological, political, and economic concerns, any one of which has enough material for many volumes in its own right. Despite these differences and the staggering scale over which they present themselves, though, it is possible to speak to general trends in the structure and development of the electric grid that hold true throughout the world.

The Division of the Grid

Generally speaking, grid infrastructure can be divided into several levels, corresponding to different voltages. At the highest voltage level is the transmission grid. Frequently operated at hundreds of kV, e.g. 69-675kV in the United States [74], the transmission grid is

responsible for moving huge quantities of power across large distances with low losses. The efficient operation of the transmission grid is financially critical, and failures that involve the transmission level can interrupt service to millions of customers [21]. As such, transmission equipment attracts large amounts of investment and the level of effort dedicated to optimizing its performance is extremely high.

Electrical substations are connected throughout the transmission level to handle the transition to the next levels of the grid: the lower-voltage sub-transmission and distribution networks. Transformers at these substations step voltages down from hundreds of kV to tens of kV, or to single-digit kV for smaller distribution feeders. All distribution feeders cover a very small geographical region compared to the transmission grid, and they are generally responsible for serving customers that number in the hundreds or thousands. Distribution feeders tend to be less uniform than the transmission grid, and can have a mix of one-, two-, and three-phase line segments across which electric loads are unevenly distributed. Monitoring equipment tends to be more scarce at the distribution level, and many of the analytical tools that can be brought to bear on transmission networks need to be adapted before they are applicable to distribution studies.

From the distribution level, service transformers again step voltage down to the level at which individual customers are most frequently served. Different combinations of transformer tapping and inter-phase connection can be used to produce different voltage levels at an end user's location, but 120V and 230V root-mean-square (RMS) from line to neutral are common values throughout the world. Terminology varies, but this last level of operation is frequently called the "low-voltage" or "secondary distribution" level. Depending on the context, the term "distribution" is sometimes extended to cover secondary distribution along with the primary distribution network to which it is connected.

The Grid's Evolution

For most of its existence, the grid has been a very centralized structure. Technological constraints and economies of scale have meant that most electric power is created by huge power plants, which interface directly to the transmission grid through large step-up transformers. From those points of connection, the transmission grid carries the power to customers that may not be located anywhere near those generating plants. This is coupled to the need to operate transmission grids at very high voltages: when the source of generation can be located hundreds of miles from the end user,[1] lower-voltage operation becomes impossible.

As technology has advanced, however, alternatives to centralized generation have started to emerge. Particularly in recent years, the increasing cost-effectiveness of smaller-scale generating technologies has driven an uptick in the number of generation sources that are connected to the distribution or secondary distribution grids, rather than at the transmission

[1] Washington State's Grand Coulee dam, the largest hydropower generator in the United States, is responsible for serving customers as far away as Arizona and New Mexico.

level [30]. These non-transmission generators are commonly referred to as distributed energy resources (DERs).

There is no universally agreed-upon definition of DER, which can take into account a huge variety of asset types. A common example is a grid-tied inverter that draws on a renewable energy source, such as photovoltaic (PV) or wind power, as its prime mover. But, the term also frequently refers to smaller-scale, non-renewable sources such as diesel generators. Grid-tied batteries are often considered DERs, which would include the batteries in electric vehicles (EVs) when they are recruited for grid services. In some cases, the term DER is even used to describe generation avoided with the strategic activation or deactivation of controllable loads, a process also known as demand response. In this dissertation, though, the term DER will only be used to describe power generation; demand response will be considered a separate concept.

DERs have a number of advantages when compared to large, remote power plants. The ability to co-locate distributed generators with the loads that they serve leads to increased reliability for end users, as there is less risk of interruption due to a contingency elsewhere on the grid. For utilities, power losses and costly upgrades to infrastructure can be minimized if the amount of power that needs to be transmitted over significant distance is reduced by increasing DER deployment. DERs can also participate effectively in a number of grid-support services, such as the maintenance of an acceptable voltage profile throughout a network. This is a function which, under the model of centralized generation, requires voltage-regulating transformers and other specialized equipment. In addition, because of the ease with which renewables can be integrated into the distributed generation paradigm, a movement of the generation balance towards DERs would favor the replacement of carbon-intensive sources of energy with cleaner alternatives.

There do exist significant challenges alongside DERs' advantages: the intermittent nature of many renewable sources, such as wind or PV, needs to be managed in order to deliver reliable service, and there is potential for adverse interactions between DERs and legacy grid equipment [69]. Economic concerns also exist, in part because highly localized and user-owned generation represents losses in incoming revenue from utility customers, a deficit that affects the financing of infrastructure maintenance.

Those difficulties will have to be faced, though, and in fact many grid operators are already grappling with their early manifestations. As of 2016, nearly one in five electric customers in the state of Hawaii had a home PV system, as did approximately one in ten single-family homes in California [54]. Though these two states owe some of that growth in PV installation to mixtures of economic and policy incentives, which vary by region, all trends indicate a broader rise in PV installations across the globe into the foreseeable future [59]. There have been many similarly encouraging reports on the potential for combined heat and power (CHP) generation [26], grid-tied batteries [75], and other power sources that are commonly connected to the grid at the distribution level. Though the current pandemic has upended predictions for the immediate future across all industries, high confidence remains in projections for a longer-term growth trend in distributed energy.

As the grid accommodates ever-increasing levels of DER penetration, the control frame-

works that govern those DERs' power dispatch have become a central concern for system operators. In Chapter 2, we will cover several broad categories of DER control strategies that have been proposed and, in some cases, implemented on real distribution feeders. PBC, a relatively recent addition to the academic literature and the focus of this dissertation, will be described in detail, as will previous work on its ability to harness DERs to achieve various objectives at the distribution level.

Microgrids and Minigrids

In many ways, microgrids and minigrids are the logical endpoints of attempts to maximize the benefits of an increasing penetration of DERs into the worldwide energy mix. The two terms are frequently used interchangeably, or with "microgrid" as an umbrella term that encompasses "minigrid," but some works draw a distinction between the two concepts.

In this dissertation, we will adopt the following definitions:

- A **microgrid** is a portion of a broader distribution-level electric grid that has the capability to disconnect and run independently from the rest of the grid, though not necessarily indefinitely. This aligns with the widely cited formal definition of a microgrid from [72]: *A group of interconnected loads and distributed energy resources within clearly defined electrical boundaries that acts as a single controllable entity with respect to the grid. A microgrid can connect and disconnect from the grid to enable it to operate in both grid-connected or island-mode.*

- A **minigrid** is a small-scale, stand-alone electrical network with no connection to a larger grid.

These concepts are similar, but they differ in what constitutes their normal operation. Always common to both definitions, though, is the ability to operate as a power island; i.e. without a connection to a broader, nation-scale electric grid. This island-mode operation of micro- and minigrids will be the focus of Chapter 4, in which the PBC strategies developed for the distribution grid are extended to isolated networks.

Microgrid Conceptual Origins and Use Cases

Formal research efforts on microgrid control strategies trace back to the early 2000s, with conceptual development by the Consortium for Electric Reliability Technology Solutions (CERTS) [39]. The CERTS concept was motivated by the drive to integrate DERs into the power system, and it focused on coordinating sets of DERs so that they could present themselves to the broader grid as a single operating unit. That unit was meant to be modular and easily extendable to accommodate new types of resources as local DER penetration increased.

To that end, CERTS developed a peer-to-peer control structure in which generators within the microgrid could determine appropriate levels of power injection without relying on

central controllers or vulnerable communication links. In many later works, the organization expanded upon their approach and described strategies for islanding, resource management, and multi-microgrid autonomous control (e.g. [58], [38]). Though microgrid research has grown over the last two decades to include participants from many academic and commercial institutions, CERTS remains a fixture in the field today.

Presently, microgrids are most commonly found on the campuses of large institutions that both generate and consume significant amounts of power. An example that often appears in literature is the Santa Rita Jail in Alameda County, CA, which includes mutliple MWs of diesel, PV, and fuel cell generation capability [27]. Another is the University of California, San Diego, which operates a 42MW microgrid [78] that integrates gas turbines, PV, and CHP systems. The U.S. military has engaged in microgrid R&D, e.g. [70], and begun developing microgrids serving Marine Corps Air Station Miramar and Parris Island. The rationale for these institutional deployments tends to be focused on additional service reliability, but economics and the improved integration of variable energy resources are also widely acknowledged benefits.

Though it could be applied to institutional networks, the work of this dissertation is most relevant to a new category of microgrid: a utility-connected community network with the ability to operate in island mode. Though this type of community microgrid is not yet common, a variety of actors have put forward different visions for potential architectures. These range from microgrids serving customers at a neighborhood level to single buildings that can form multi-source microgrids internally [49]. As fundamental as the physical layout of these community microgrids is their establishment as a new class of legal entity, with accompanying social and economic concerns. The daunting task of defining these aspects of a model community microgrid is being undertaken by numerous experts, one exemplar effort being the Oakland pilot installation of UC Berkeley's large, multidisciplinary EcoBlocks project [3].

Within the emerging space of community microgrids, PBC may be especially well-suited for a subtype that could become possible in the medium-term future: a microgrid developed at the primary distribution level through a grassroots, bottom-up process. This type of development would be very different from the paradigm that exists now, in which microgrids are carefully designed and centrally planned.

In a bottom-up development strategy, microgrids would be created organically. As residential DERs continue their penetration into the grid, there will likely be portions of distribution feeders that find themselves with sufficient generating capacity to serve local loads for some portion of the day, even in the absence of connection to the broader grid. In that case, grid operators could offer incentive programs that allow the owners of those residential DERs to opt into the provision of grid services, which could include ceding control of their DERs to the operator under pre-specified circumstances. With the proper communications infrastructure in place and the switching capability to sectionalize self-sufficient portions of distribution feeders, those residential DERs could be used to form ad-hoc microgrids capable of harnessing the benefits of aggregated residential generation.

Off-Grid Power with Minigrids

Minigrids are used to serve isolated areas that can't be easily connected to an established power grid. They vary in size and sophistication, and the specifics of their design are very dependent on their application. The financials of minigrids are very different from those of larger grids, as they lack economies of scale, and they face physical challenges in serving loads without a large generation buffer. Because there are stability advantages to having the inertia of rotating generators within a power system, many currently operational AC minigrids are hybrid, i.e. they draw their power from a mixture of inverter-interfaced renewables and spinning, fossil-fuel generators. All-inverter minigrids are possible as well, though, and inertia-free minigrid control is an active area of research [11].

Many minigrids are deployed by industrial operators. Resource extractors, particularly mining operations, have begun exploring the potential of minigrids to cheapen their energy consumption and decrease their environmental impact [44].

Community-scale minigrids make up another category of application. By necessity, island communities have been early minigrid adopters [13]. Remote continental areas, too, have seen deployments; a number of towns in Alaska have drawn their power from minigrids for many years and are actively involved in research and development with universities [29]. The most technically advanced community minigrid deployments have, in large part, been enabled by grants and other governmental funding. Fully commercial projects are less common, due to high up-front costs and the generally low population density of the customers to be served by the finished product.

Minigrids are also frequently considered in electrification initiatives for rural communities in the developing world. These undertakings address a critical need. As of 2018, the number of people worldwide who lacked access to electricity stood at 789 million [73]. Though there has been progress made towards the goal of universal electric access in the past decade, it will remain a significant humanitarian issue in the years to come: according to the World Bank, the estimated effects of projecting current policies into the future will leave 650 million people still lacking access to electricity in 2030 [46].

It should be mentioned that the majority of currently operating microgrids and minigrids, particularly those in the developing world, are run at or below secondary distribution voltage levels and would not be candidates for the PBC discussed in this dissertation. However, as the sizes of future independent, DER-based networks grow, the medium-voltage level and centralized control strategies are likely to become more relevant to those projects.

1.2 Three-Phase Coordination

PBC could be an effective strategy for dispatching DERs across both minigrids and community microgrids in islanded operation. One important objective that it could help to achieve is the coordination of single-phase DERs to present a balanced, three-phase voltage profile across an islanded network.

Modern AC electric grids are built on the principle of three-phase power generation. Under this system, power is transferred using a minimum of three wires, which carry sinusoidal currents that are nominally offset from one another by 120° in phase. In the United States, these three waveforms are usually labeled *A*, *B*, and *C*.

On many primary distribution feeders in the U.S., three-phase power is delivered using wiring configurations that vary by branch. So-called three-phase-four-wire (3PH4W) branches are used to serve Wye-type loads, which are connected between the live conductor of phase A, B, or C and a neutral wire. Where the loads are of Delta type, meaning that they are connected between the live conductors of adjacent phases, they can be served with three wires. Branches serving fewer customers will sometimes use only two wires, a single phase and neutral. Individual customers are often connected to the low-voltage side of a single-phase service transformer, regardless of their primary distribution branch type.

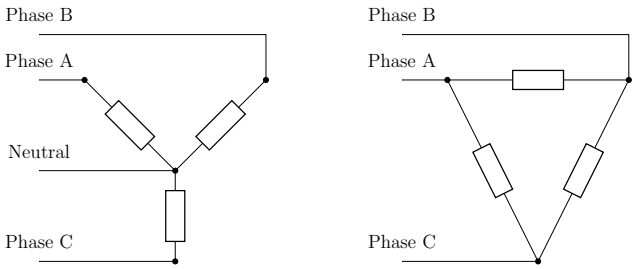

Figure 1.1: Generic Wye- (left) and Delta-connected (right) loads

Three-phase power delivery has several advantages when compared to a two-wire, single-phase alternative. The first is that three times the amount of power can be delivered through the circuit while using proportionally less conductive material, because the three live conductors of a 3PH4W circuit are able to share a single neutral wire. In addition, that neutral wire can have a smaller cross section than would be required on a single-phase circuit, as the relative 120° shifting of the currents on the live wires means that no current will flow on the neutral under balanced operation.

Another advantage to three-phase power delivery is that it results in more consistent power delivered to a load over time. This is not an issue for most common electric loads; AC voltage and current on a single-phase line is standardized at a nominal value of 50Hz or 60Hz throughout the world, and the power transferred on a single-phase circuit fluctuates at twice that frequency, which is generally too fast to affect connected devices. However, large motors, which are built to take advantage of three-phase power, are an important exception.

Smaller, residential DERs of the type that might make up an ad-hoc community microgrid generally provide single-phase power, and they are meant to connect at a single-phase customer's service panel. Those residential DERs serve the single-phase electric load on site and, where agreements with the local utility are in place, pass excess generation to the grid on the phase to which they are connected. Single-phase, grid-tied DERs synchronize them-

selves to the local grid voltage, aligning the phase angles of their voltage waveforms using those of either phase A, B, or C.

In Chapter 4, we explore the coordination of single-phase DERs to provide three-phase power in the case where a strong grid reference does not exist. This situation could arise on a portion of the distribution grid that has been islanded during a contingency, or it could be the permanent operating state of a four-wire minigrid. In either of those scenarios, the ability to create a three-phase electric grid from locally available single-phase resources could allow motors and other sensitive loads to operate on network configurations that would otherwise only allow for single-phase loads.

1.3 Battery-Coupled DERs and Complex Power

The DERs that provide power for microgrids and minigrids are not of a uniform type. They vary in both generation capacity and ownership model, from relatively large commercial or utility generators connected on the primary distribution grid to the small, customer-sited systems described in the previous section. An emerging subtype of the latter is the residential-scale battery.

Residential batteries are most commonly connected to the grid as part of solar-plus-storage systems, in which they are used to collect excess energy from a PV installation for the purpose of energy arbitrage or backup power. In these systems, the PV array and battery can be connected in several ways: through a charge controller on the DC side of a bi-directional converter, through the AC sides of individual converters, or through a specially designed hybrid converter. As a relatively new category of residential storage, some companies have begun offering "AC batteries" pre-packaged with dedicated converters. In the United States and Europe, Tesla's *Powerwall 2* is a highly recognizable AC battery offering, and the launch of ZOLA Electric's *Infinity* represented the first AC battery developed for African markets.

For the control strategies developed in this dissertation, we will take battery-based DERs as our archetypal generator. This allows us to assume that, at least over short operational periods, the intermittency of wind- or solar-based prime movers will not affect the ability of our DERs to provide consistent power. We will also assume that the DERs' converters are capable of "four-quadrant" operation, meaning they can be commanded to dispatch both active and reactive power as well as charging or discharging their battery.

The distinction between active (often called "real") and reactive power is an important construct that describes the degree to which voltage and current waveforms are aligned in phase, and thus how effective they are at performing useful work over time [43]. Because of the centrality of these active and reactive power concepts to the work of this dissertation, we provide a brief derivation here. As an illustrative example, consider a single-phase electrical load, modeled as a parallel connection of a resistor and an inductor, connected to a stiff, sinusoidal voltage source.

This section loosely follows the presentation of [25]. We begin by defining the fundamental components of voltage at the load's point of connection and the current flowing through the

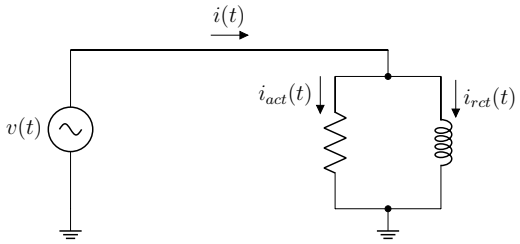

Figure 1.2: A representative load connected to a stiff power grid

load:

$$v(t) = V_{max}\cos(\omega t) \qquad i(t) = I_{max}\cos(\omega t - \theta) \tag{1.1}$$

where θ is the phase delay of the current waveform relative to the the voltage and ω is the system frequency. Neglecting harmonics, the instantaneous power $p(t)$ drawn by the load can then be defined and its expression rearranged:

$$\begin{aligned} p(t) &= v(t)i(t) \\ &= V_{max}\cos(\omega t) I_{max}\cos(\omega t - \theta) \\ &= V_{max}I_{max}[\cos(\theta)\cos^2(\omega t) + \sin(\theta)\sin(\omega t)\cos(\omega t)] \end{aligned} \tag{1.2}$$

We can see that $p(t)$ is the sum of two components. The first has an average value of $\frac{1}{2}V_{max}I_{max}\cos(\theta)$ over an integer number of cycles, and the second averages to zero.

We then introduce two new definitions, i_{act} and i_{rct}:

$$i_{act}(t) = I_{max}\cos(\theta)\cos(\omega t) \qquad i_{rct}(t) = I_{max}\sin(\theta)\sin(\omega t) \tag{1.3}$$

that allow us to rewrite Equation (1.2) as

$$p(t) = v(t)[i_{act}(t) + i_{rct}(t)] \tag{1.4}$$

i_{act} captures the portion of the current in phase with the voltage and i_{rct} describes the component that is 90° out of phase.

We define the active power P as the average value of $v(t)i_{act}(t)$. The reactive power Q is defined as the maximum of $v(t)i_{rct}(t)$ over a cycle. Rewriting the amplitude coefficients for voltage and current in terms of their RMS values gives us P and Q in their most commonly presented forms:

$$P = V_{rms}I_{rms}\cos\theta \qquad Q = V_{rms}I_{rms}\sin\theta \tag{1.5}$$

The Pythagorean sum of P and Q is known as the "apparent" power and denoted by S

$$\begin{aligned} S &= P + jQ \\ &= V_{rms}I_{rms}(\cos\theta + j\sin\theta) \end{aligned} \tag{1.6}$$

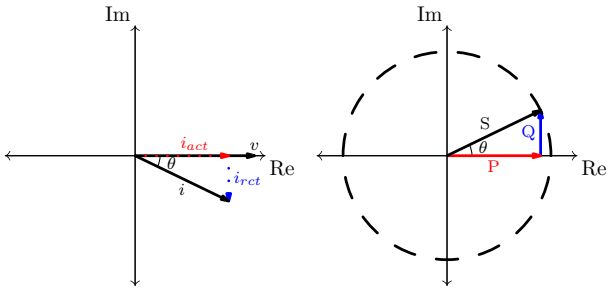

Figure 1.3: The graphical relationships between voltage and current (left) and active, reactive, and apparent power (right)

Active, reactive, and apparent power are useful quantities in power systems analysis. Active power, measured in Watts, is what we most naturally think of as the power delivered by an electrical grid. It represents the energy over time that carries out useful work in devices connected to the network.

Reactive power, measured in Volt-Amperes reactive (VARs or VArs), describes the power that is passed between generators and the inductive or capacitive elements of electric load. Averaged over a cycle, reactive power does no useful work. However, it represents real circulating currents in power lines that, among other effects, contribute to resistive losses.

Apparent power, measured in Volt-Amperes (VA), is directly tied to the maximum instantaneous power over a cycle, and is thus a convenient quantity for expressing constraints on the capacity of generators or power transmission elements.

As stated above, we will assume in this work that the only restriction on our controllable DERs is a physical bound on apparent power. In practice, many commercially available inverter or AC battery options are restricted by design or by local regulation from operating in true, unrestricted four-quadrant mode. But, as the grid evolves and DERs take an increasingly active role in its operation, both design motivations and regulatory environments are likely to change.

1.4 Structure of this Dissertation

PBC is a recently developed strategy, and it differs markedly from other approaches to the control of DERs. Chapter 2 will be entirely devoted to a description of PBC as it was conceptualized by our research group for DER control on distribution feeders. It will also introduce the Linearized Unbalanced Power Flow Model (LUPFM), developed by [62], which provides a means for implementing the supervisory control layer that forms an integral part of the PBC structure. It is the development of this supervisory layer that will be the focus of the work to follow.

Though the primary subject of this dissertation is the application of PBC to power systems operating in island mode, it was first necessary to develop a workable implementation of PBC's supervisory controller. A particular challenge was addressing the low tolerance for error in the supervisory control layer's calculations, which presents a significant difficulty for the practical use of PBC at the medium-voltage level. In Chapter 3, we describe an adaptation of an iterative solution method from [62], which allows us to obtain near-exact solutions to nonlinear power flow equations while making use of a linearized network model. We further adjust the solution method to incorporate tap-changing transformers, with the end result being a supervisory controller implementation that can generate phasor targets of sufficient accuracy to enable PBC on realistic distribution feeders.

We then present the island-mode extension of PBC. This involves the development of a slack bus placement strategy, which will be discussed in the first part of Chapter 4, and the correction of a particular failure mode of our controller. Using both of these inputs, we develop an end-to-end implementation of the supervisory control layer and demonstrate it in simulation for a network with several levels of DER penetration.

Finally, we return to the practical challenges of implementing PBC, discussing an issue that will be of significance for many PMU-based[2] monitoring and control applications: the measurement error induced by instrument transformers. Chapter 5 presents an empirical view of the change in that induced error over time, which is not generally accounted for in measurement models, and estimates its potential effects on PBC.

Each of Chapters 2, 3, 4, and 5 will introduce their subject and provide a review of relevant literature. Where applicable, they will present simulation or other results, and describe their significance in the context of this dissertation. Final conclusions and a description of possible directions for future research will be included as Chapter 6.

[2]PMUs, which are used to carry out synchrophasor measurements, will be discussed in Chapter 2

Chapter 2

PBC and Research Preliminaries

This chapter covers our definition of phasor-based control (PBC) and the research that laid groundwork for the remainder of the dissertation. In Section 2.1, it introduces the PBC concept as it was developed for non-islanded distribution systems. In Section 2.2, it briefly describes the optimal power flow (OPF) problem and its role within PBC. Section 2.2 also reproduces a derivation of the Linear Unbalanced Power Flow Model (LUPFM), which provides constraints for a quadratic program (QP) used by our proposed implementation of PBC.

2.1 Phasor-Based Control

Recent advancements in the sensing technologies that can be deployed on the electric grid have opened up new possibilities in the control of DERs. Categories of information that were previously only economically available on transmission networks can now be harnessed in the operations of distribution feeders and micro/minigrids. This expansion of sensing capabilities has enabled the consideration of fundamentally new strategies for dispatching DERs in a way that not only contributes to healthy grid operations, but that allows operators to pursue objectives that have not previously been possible at the distribution level.

PBC is one of these proposed strategies [76]. It was developed as a hierarchical control framework; at the topmost level, a centralized supervisory controller solves an OPF problem. That solution is then used to generate a set of phasor reference values, which are provided to a set of local controllers throughout an electrical network. Each of those local controllers has jurisdiction over DERs installed on a subsection of the network and access to a voltage phasor measurement within the same area. The task of the local controllers is to dispatch their DERs to meet the supervisory controller's specified phasor reference and to maintain that value in the face of load steps or other changes in conditions on the feeder.

Phasors and PMUs

The Definition of a Phasor

PBC is unique in that it makes direct operational use of voltage phasors, which are structures that can be used in simplifying AC circuit analysis. In steady-state, phasors compress information about a sinusoidal quantity $x(t) = X_{max}\cos(\omega t + \theta)$ into a pair of values $X_{max}\angle\theta$ or $X_{rms}\angle\theta$, where $X_{rms} = \frac{X_{max}}{\sqrt{2}}$.

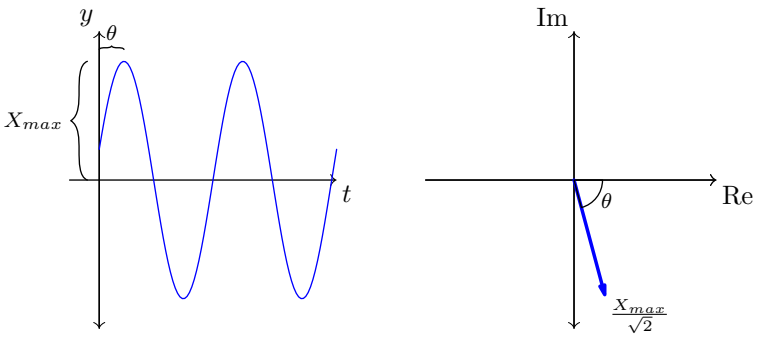

Figure 2.1: A phasor representation of a sinusoidal waveform

Intuitively, a phasor can be considered a vector that rotates counterclockwise in the complex plane with frequency ω. At any point in time, the instantaneous value of the underlying sinusoid $x(t)$ will be given by the projection of the phasor onto an axis. A phasor diagram, of which Figure 2.1 is an example, displays the phasors of all relevant quantities at a snapshot in time, usually the $t = 0$ reference.

There are two degrees of freedom in the definition of a phasor. In addition to the decision to specify magnitude in absolute or RMS value, the mapping from a phasor to an instantaneous value of $x(t)$ can be defined as the projection of $X_{max}e^{j(\omega t+\theta)}$ onto either the real or imaginary axis. This dissertation will use the convention of [56], which represents a steady-state phasor's magnitude as its RMS value and the instantaneous value of the underlying sinusoid as the phasor's projection on the real axis multiplied by $\sqrt{2}$.

Implied in the definition is that the angle θ is provided with respect to a time reference common to all phasor measurements on a network. It is also implied that a single steady-state frequency ω is shared by all phasor quantities. This latter characteristic means that phasors must be defined independently for each of the harmonic components of a physical quantity such as voltage or current. In power flow studies, those physical quantities are often approximated as consisting only of the fundamental component of the waveform; other tools are used to treat harmonic distortion if desired.

Phasors are exceptionally convenient for the steady-state treatment of AC power systems, but they can also be extended as a tool for dynamic analysis. While this is not explored in depth as part of this dissertation, interested readers are encouraged to refer to [61] and [67] for more information.

Phasor Measurement Units

As a concept, phasors have existed for over a century, but their practical use in physical measurement is a more recent development. Phasor Measurement Units (PMUs) were first built at Virginia Tech in the 1980s, where they were designed to efficiently convert sampled data to phasor quantities using algorithms developed for computerized relaying [55]. For clock synchronization across widely separated units, previously a fundamental hurdle in development, these prototype PMUs made use of common timestamps provided by the recently launched GPS network. In the early years of incomplete satellite deployment, PMUs required GPS receivers with expensive internal crystal clocks to maintain synchronization, but relatively cheap receivers are now capable of providing timing that is accurate to within a microsecond, which translates to an angular error of about 0.02° for a modern PMU installed on a 60Hz system [57].

PMUs were first offered commercially in the early 1990s. They quickly became integral to grid operations, where they are used today in Wide Area Monitoring Systems (WAMSs). Well into the last decade, though, PMUs were near-exclusively deployed on the transmission network. In large part, this was due to economics: because of the much greater number of customers dependent on the transmission grid, transmission monitoring is a more critical task that attracts more investment than does the monitoring of any particular distribution feeder.

There are also technical complications involved in deploying PMUs at the distribution level, in particular the need for very high measurement precision. The lower amounts of power flowing through distribution networks translate into smaller angular separations between voltage waveforms at adjacent nodes, meaning that higher-fidelity phasor data must be captured in order to extract useful distribution-level information. In recent years, PMUs designed specifically for the distribution grid have achieved this level of precision and become available, both embedded in system components and as commercial products in their own right. An early example of the technology, the μPMU, is capable of measuring phasor angles with a typical accuracy of 3 millidegrees [45].

There are obstacles yet to be overcome in accessing the full potential of distribution-level PMUs in operational settings. One of those obstacles is the systematic error introduced by instrument transformers, which will be examined in detail in Chapter 5. However, despite challenges, the potential distribution-level applications for PMUs are exciting and varied; a survey can be found in [77]. Most relevant for this dissertation, distribution PMU measurements can be integrated into control schemes for dispatch of DERs.

The PBC Strategy

In the introduction, we briefly touched on the need to control DERs in a way that benefits, rather than hinders, the operations of the distribution grid. This requires an awareness of the potential for local congestion on the network, violations of allowable voltage limits, and reversal of the expected direction of power flow on the feeder [14], any of which will influence

the amount of power that should be dispatched by a DER at a given time. The difficulty of operating at the distribution level is compounded, as well, on feeders where where there is a lack of reliable model or measurement information [68].

A vast number of DER control strategies have been proposed by actors in both industry and academia. These have ranged from centralized methods, as in [36], to the fully local, such as the Volt/VAR control scheme [66] that is required functionality for advanced inverters under California's Rule 21 [42]. Between those extremes, decentralized and distributed control techniques abound; a review of them can be found in [6].

Other research has focused on market-driven approaches to DER control, in which pricing and operational constraints combine to influence power contributions in ways that recall ISO strategies for managing the transmission grid [20]. Many of these economic strategies rely on Locational Marginal Pricing (LMP), in which the price of electricity varies between nodes on a grid in order to incentivize appropriate behavior from DER operators and other actors. LMPs could be used in a very responsive way at the distribution level with high spatial granularity [40], but there are potential drawbacks to creating highly localized fluctuations in electricity prices that end users may find frustrating.

PBC provides an alternative to traditional approaches to DER dispatch by recasting power setpoints in terms of voltage phasor targets. It is built on the idea that a great deal of information about the present operating state and stability of an electric network can be conveniently expressed in terms of phasors, and that important operating constraints can be encoded into the phasor profile of a distribution network.

As mentioned, PBC is a hierarchical, layered control scheme. At the top layer, a centralized supervisory controller, referred to as the S-PBC, solves an OPF over an electrical network with respect to an arbitrary objective. PBC is meant to be a plug-and-play control framework, so it does not prescribe the solution technique to be used for the S-PBC. The only constraint on the supervisory control scheme is that it be capable of returning complex-valued phasors as nodal voltage targets.

Once the S-PBC has solved its OPF, the phasor values returned by the solution will be communicated to the next-lowest level of control. In principle, that could be the lowest-level control layer, which is directly responsible for actuating DERs, or it could be an intermediary controller responsible for operations within a portion of the higher-level S-PBC's network. In either case, it is important to point out that the higher-level S-PBC does not need a detailed model of the portion of the network governed by the lower-level controllers. It only requires a few basic pieces of information about that sub-network, such as the total active and reactive power that can be requested from the lower-level controller. In effect, the higher-level S-PBC's model will end at the point that the next-lower-level controller's model begins. To the higher-level controller, the lower-level controller's network will appear as a node from which a given amount of power can be commanded.

At the lowest level of operation sits a local controller, which we designate the L-PBC. Though the L-PBC pursues a similar goal to any intermediate layers that may be present, i.e. the tracking of a phasor target provided by the next-highest control level, the L-PBC structurally differs in that it is directly responsible for individual DERs. As was the case for

the S-PBC, the PBC framework is meant to be agnostic to the specifics of the L-PBC. But, where the S-PBC is meant to solve high-level optimization problems, the L-PBC is tasked with the real-time maintenance of a phasor target, which generally lends itself to feedback controllers or other fast-acting strategies. In searching for suitable L-PBC controllers, our research group and collaborators have explored a diverse set of approaches, including PI, extremum-seeking [63], and Retrospective Cost-Adaptive Control [33].

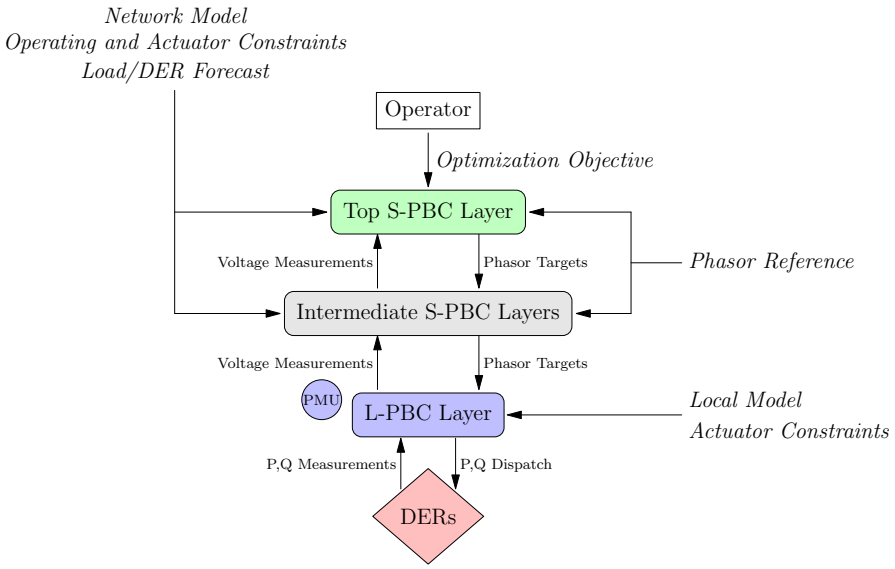

Figure 2.2: The control layers of the PBC strategy

The division between the roles of the S-PBC and L-PBC is at the core of the PBC strategy. The S-PBC layer is meant to ensure that an optimal steady-state operating condition is provided to lower-level controllers in the form of phasor targets to be tracked. These targets are not meant to encode instructions that are critical to the safety or stability of the system, and so they can be updated at a slow rate compared to that of the L-PBC's operation. This gives the S-PBC the freedom to solve larger, system-wide optimization problems. The L-PBC, on the other hand, is responsible for responding immediately to maintain the S-PBC's phasor targets in the face of DER prime mover fluctuations, load steps, or other similar disturbances. If conditions on the network change between S-PBC control updates, the expectation is that the local controllers will alter their DER dispatch in an attempt to maintain the previously assigned voltage phasor values. Those values would almost certainly no longer be optimal in the sense of minimizing a system-wide objective, but they would represent an operating condition that is known not to be critically harmful. The feeder would remain in that condition until a state estimator updated the S-PBC's understanding of the network and newly optimized phasor targets could be assigned.

As illustrative examples of PBC operation, two toy problems presented in [76] are briefly reproduced here. These problems assume a two-layer PBC made up of a supervisory and

a local control level, without any intermediate controller stages. This two-level controller assumption will be the default throughout this dissertation, as it allows for the exploration of the unique characteristics of PBC without introducing the unnecessary complication of intermediate control layers.

The first of the toy problems demonstrates the ability of PBC to localize its DERs' responses to changing conditions on a distribution feeder. Our notional feeder will be the radial example of Figure 2.3. In the initial steady state, the DER has been tasked with driving its point of connection to the phasor value $V_{DER}\angle\theta_{DER}$, with θ_{DER} measured relative to the substation's voltage phase angle and magnitude V_{DER} measured relative to neutral. The DER will then act to cancel out any changes to that voltage caused by changes in loading on the feeder, which will translate to adjusting its dispatch to assume any new local power demand or to reducing its output to reflect load that has been disconnected.

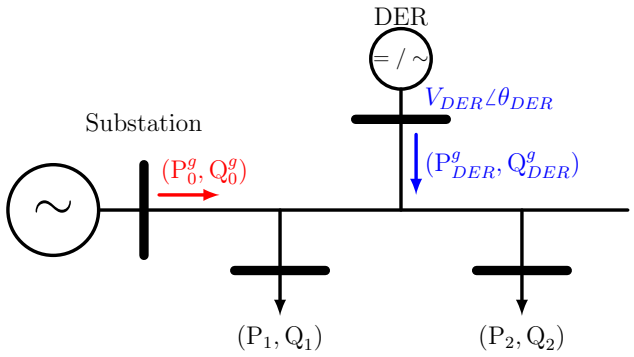

Figure 2.3: Toy problem 1: A demonstration of localized DER response under PBC

While a similar re-dispatch of DERs could be accomplished using measurement of incoming power P_0 from the transmission system, there are important cases where that would be problematic. Consider a situation in which the load P_1 decreases by ΔP and the load P_2 increases by the same amount. Neglecting losses, that change would not appear in P^g_0, but the additional distance over which the quantity ΔP must now travel could result in a voltage violation at the end of the feeder. In contrast, if the DER is operating under PBC, it will immediately react to its change in local voltage phasor by increasing its production to assume the additional load.

A second use case in which PBC is particularly effective is DER reassignment following loss of line capacity. Consider the situation diagrammed in Figure 2.4, where two lines of equal impedance, Line $Ln\alpha$ and Line $Ln\beta$, carry power from Bus 1 to Bus 2. Each of the lines carries power of magnitude P_{12}, Q_{12}. If one of those two lines is lost, the remaining line will now be responsible for carrying all power transferred between the two nodes; in a worst-case scenario, that would mean the line must carry double its thermally rated current until the DER dispatch is updated.

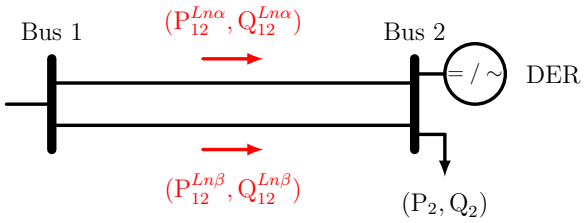

Figure 2.4: Toy problem 2: PBC's response to a broken conductor

Under PBC, DERs immediately take action to rectify the situation. The loss of a parallel current path causes an increase in the effective resistance between Bus 1 and Bus 2. The transfer of $2P_{12}$ and $2Q_{12}$, then, causes a larger voltage change across the line and drives the DER at Bus 2 to increase its production to counter it. Assuming sufficient DER capacity is available at Bus 2 to restore the phasor to its original value, the power across the surviving conductor would be reduced to its original P_{12}, Q_{12}, with the rest of the power demanded by load at Bus 2 being provided locally.

Even in the case that Bus 2 does not have sufficient DERs to serve the entire load, in attempting to restore the original phasor value at Bus 2 we can expect that the correct local control action will be taken, with DER dispatched to serve as much of the load as is possible given available capacity. The overload condition on the connecting line will then be alleviated to maximum extent, moving the system to a safer operating state. The system will then stay in that state until the S-PBC is able to recalculate the feeder's OPF and re-assign phasor targets.

The unique properties of PBC that allow it to respond correctly over a wide set of possible contingencies make it a promising control technique. However, it is a recent development, and there exist barriers to its practical implementation that will need to be addressed in future research. A primary challenge that must be overcome is the need for extreme precision in specifying feasible phasor targets at the S-PBC level. In Chapter 3, we will demonstrate an ability to achieve that precision in theory, but the realistically achievable accuracy of available modeling information at the distribution level remains an unsolved concern. Despite this hurdle, aspects of PBC have been demonstrated experimentally under the ENERGISE project, with results to be published in [10].

PBC Applications on Islanded Networks

When applied to minigrids or to microgrids operating in island mode, there are two operational goals that PBC is especially well-suited to achieve. Both have been previously explored in simulation on non-islanded distribution systems in [62]. Chapters 3 and 4 revisit those goals in simulation, with Chapter 4 demonstrating their extension to islanded networks.

The first goal was introduced as an idea in Section 1.2. PBC can be used to recruit DERs to create balanced, three-phase voltages across a feeder in the absence of a connection

to a stiff transmission grid. PBC is able to explicitly command that voltage balance in the communication of phasor targets between S-PBC and L-PBC, assigning voltage values that maintain 120° of separation between grid phases to the extent possible given the demands on the system.

The second PBC objective we will consider is "phasor matching." In a phasor matching use case, DERs are dispatched so as to drive the voltage at a selected node $\mathbf{V}_n$ to a specific, externally provided phasor value $\mathbf{V}_{ext}$. The goal of the OPF solver is to generate a feasible set of voltage phasor targets that can be sent to controllers at each node in the network. That target set will include a value for $\mathbf{V}_n$ that is as near as possible to $\mathbf{V}_{ext}$.

As a clarifying point, this phasor matching objective is distinct from the phasor tracking action of controllers at the L-PBC level. As part of the definition of the PBC framework, each L-PBC will be constantly actuating individual DERs to maintain a real-time voltage phasor value provided to it by the S-PBC. That S-PBC target, though, could be issued with the intent of achieving any network-wide objective, such as minimizing losses or total generation cost. When the feeder-level objective is described as "phasor matching," it refers to a case in which the S-PBC itself has been tasked with meeting a provided phasor-defined voltage reference, and that its objective is to recruit L-PBCs across its network to match that reference at a specified node.

This ability to match voltage at a given node to a specified phasor value would be useful in grid reconfiguration, and as such is more applicable to microgrids than to stand-alone minigrids. As mentioned in Section 1.2, if a contingency or maintenance requirement created the need for disconnecting a distribution-level microgrid from the transmission network, the correct control strategy would allow that microgrid to continue to run in isolation, provided that it had sufficient DER capacity over the relevant time period. Once the period of abnormal grid operation passed and it became desirable to reconnect the feeder to the broader grid, phasor matching would allow for the nodal voltage at the open tie switch to be set equal to the value on the switch's opposite side. PBC and the LUPFM are uniquely suited to this application because of their ability to explicitly match not only voltage magnitude but phase angle as well, reducing the chance of arcing due to an out-of-phase reclosure.

2.2 S-PBC and the Linear Unbalanced Power Flow Model

Optimal Power Flow

A primary task of the S-PBC is solving for optimal power flow on its feeder. As a general class of problem, OPF refers to the search for the best possible actuation of controllable equipment on an electrical network, where "best" is defined on a situational basis. It is an incredibly broad body of work, and the modern form of the problem traces its roots to the 1960s. Since that time, it has been approached through a wide array of formulations, many tailored to specific use cases.

The "full" problem, referred to as ACOPF in [28], is well-defined but difficult: the physics of power flow ensure that the system's governing equations give rise to a nonlinear and nonconvex optimization problem, equipment that operates in discrete steps imposes integer constraints on the system, and security concerns require that solutions to OPF be robust to possible losses of conducting lines or generating equipment. To this day, there is no single solution method that is agreed to be sufficiently fast and robust for universal adoption.

A truly optimal methodology for quickly solving the OPF problem would have incredible implications for power systems. Solving OPF is a core piece of the work done by Independent System Operators (ISOs), the agencies tasked with overseeing large regional transmission grids. As a general practice, ISOs must solve OPF problems on many time horizons, up to and including near-real-time[1]. It is estimated that even a 5% efficiency improvement in the power flow solution methods used by ISOs would result in $6 billion in savings in the U.S. and $26 billion worldwide. As such, OPF is an extremely active area of research focused on developing approximations, decompositions, and other approaches that can be integrated into more accurate and efficient solvers.

Much of the recent research into OPF solutions has focused on convex relaxation techniques, in which the solution space defined by system constraints is expanded to include a superset that allows for the application of second-order cone programming (SOCP) or semi-definite programming (SDP) techniques. If the chosen solver then returns an optimal set of variables that satisfies the original set of power flow constraints, that system state is a global optimum for the original problem. A description of some conditions under which convex solvers return exact solutions is given in [41], and [62] has an excellent survey of many recent works in SDP relaxation.

While many of these methods show promise and achieve success under a variety of conditions, there are outstanding difficulties in practical implementation, particularly on unbalanced, multi-phase distribution networks. As a result, we instead look to quadratic programming (QP) for implementing the S-PBC. To construct our QP, we use a linearized approximation of the power flow relations, the LUPFM, as our network model. This model defines our problem's constraints, and it allows for arbitrary quadratic objectives to be defined in terms of voltage phasors.

The LUPFM is based on the "DistFlow" equations, which were originally introduced in 1989 [7]. The DistFlow equations established relationships between the active and reactive power flowing through conducting lines and the squared voltage magnitudes at each of a network's nodes. A simplified approximation of the DistFlow equations, referred to as either "Simplified DistFlow" or "LinDistFlow," was derived from the original equations by neglecting several quadratic terms and linearizing all relationships [8]. The Simplified DistFlow equations were later generalized from their original, single-phase form to unbalanced,

[1]The FERC report [28] written by Mary B. Cain, Richard P. O'Neill, and Anya Castillo writes "This problem must be solved weekly in 8 hours, daily in 2 hours, hourly in 15 minutes, each five minutes in 1 minute and for self-healing post-contingency in 30 seconds."

multiphase networks [23].

Building on this three-phase, linear model, the work of [62] developed the LUPFM described in this section. It introduced an additional relationship between voltage phasor angles and power flows into a model that was previously only capable of specifying voltage magnitudes. This was a critical development that allowed for the treatment of a phasor in its entirety within the structure of the model. The LUPFM is presented here with minimal derivation; interested readers should refer to [62] for a more complete treatment.

Model Definition

The LUPFM equations are linear with respect to the variables $\mathbf{E}_n$, $\mathbf{\Theta}_n$, $\mathbf{p}_n^g$, $\mathbf{q}_n^g$, $\mathbf{P}_{nm}$, and $\mathbf{Q}_{nm}$. To establish a fully linear set of equations, nonlinear quantities based on current magnitudes, such as line power losses, are replaced with estimated constant values. In the equations presented in this section, those estimated values will be denoted by the superscript *est*.

LUPFM Nomenclature

$\mathcal{N}$ Set of all nodes on the distribution feeder

$\mathcal{E}$ Set of all connecting lines on the distribution feeder

$\mathcal{P}_n; \mathcal{P}^k; \mathcal{P}^\phi$ Subset of phases {A,B,C} present on node n; set of nodes with k phases present; set of nodes with phase $\phi \in \{A, B, C\}$ present

V_n^ϕ Complex-valued voltage phasor on phase ϕ of node n, magnitude given in per-unit RMS

E_n^ϕ Squared per-unit RMS voltage magnitude on phase ϕ of node n.

θ_n^ϕ Voltage phasor angle on phase ϕ of node n

$p_n^{d\phi}$ Per-unit active power demanded by the load on phase ϕ of node n

$q_n^{d\phi}$ Per-unit reactive power demanded by the load on phase ϕ of node n

$p_n^{g\phi}$ Per-unit active power generated by DER on phase ϕ of node n

$q_n^{g\phi}$ Per-unit reactive power generated by DER on phase ϕ of node n

$s_n^{R\phi}$ Per-unit apparent power rating of DER on phase ϕ of node n

I_{nm}^ϕ Complex-valued current phasor on phase ϕ of line nm, magnitude in per-unit RMS

P_{nm}^ϕ Per-unit active power flowing from node n to node m on phase ϕ of line nm

Q_{nm}^ϕ Per-unit reactive power flowing from node n to node m on phase ϕ of line nm

P^{ϕ}_{Lnm} Per-unit active power loss on phase ϕ of line nm

Q^{ϕ}_{Lnm} Per-unit reactive power loss on phase ϕ of line nm

$R^{\phi\psi}_{nm}$ The resistive part of element ϕ, ψ of the 3x3 impedance matrix of the line connecting nodes n and m

$X^{\phi\psi}_{nm}$ The reactive part of element ϕ, ψ of the 3x3 impedance matrix of the line connecting nodes n and m

$Z^{\phi\psi}_{nm}$ The complex-valued impedance $R^{\phi\psi}_{nm} + jX^{\phi\psi}_{nm}$

$\mathbf{\Lambda}_n, \mathbf{\Lambda}_{nm}$ If $\Lambda^{\phi}_{n,nm}$ is a per-phase quantity such as voltage, current, or power, bold variables denote 3x1 vectors: $\begin{bmatrix}\Lambda^a_n & \Lambda^b_n & \Lambda^c_n\end{bmatrix}^T$
If $\Lambda^{\phi\psi}_{n,nm}$ is a cross-phase characteristic such as resistance, reactance, or impedance, bold variables are 3x3 matrices with elements $\Lambda^{\phi\psi}_{nm}$, where $\phi, \psi \in \mathcal{P}_{nm}$

c_{Kn} Coefficients associated with the ZIP model components of a load at node n. $K \in \{Z, I, P\}$ and $c_{Zn} + c_{In} + c_{Pn} = 1$.

$\boldsymbol{\alpha} \circ \boldsymbol{\beta}$ The Hadamard (element-wise) product of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$

$|\Lambda^{\phi}|, |\mathbf{\Lambda}|$, The magnitude of element Λ^{ϕ}, the 3x1 vector of magnitudes of the elements of $\mathbf{\Lambda}$

Real and reactive power balance

Kirchhoff's Current Law is enforced at each node in the system with Equation (2.1).
$\forall n \in \mathcal{N}$:

$$\begin{aligned}\mathbf{p}^d_n - \mathbf{p}^g_n &= \sum_{ln\in\mathcal{E}} \mathbf{P}_{ln} - \sum_{nm\in\mathcal{E}} \mathbf{P}_{nm} - \sum_{nm\in\mathcal{E}} \mathbf{P}^{est}_{Lnm} \\ \mathbf{q}^d_n - \mathbf{q}^g_n &= \sum_{ln\in\mathcal{E}} \mathbf{Q}_{ln} - \sum_{nm\in\mathcal{E}} \mathbf{Q}_{nm} - \sum_{nm\in\mathcal{E}} \mathbf{Q}^{est}_{Lnm}\end{aligned} \tag{2.1}$$

where $\mathbf{P}^{est}_{Lnm}$ and $\mathbf{Q}^{est}_{Lnm}$ are loss terms, defined as

$$\begin{aligned}\mathbf{P}^{est}_{Lnm} &= \mathbf{Re}((\mathbf{Z}_{nm}\mathbf{I}^{est}_{nm}) \circ \mathbf{I}^{est*}_{nm}) \\ \mathbf{Q}^{est}_{Lnm} &= \mathbf{Im}((\mathbf{Z}_{nm}\mathbf{I}^{est}_{nm}) \circ \mathbf{I}^{est*}_{nm})\end{aligned}$$

In this expression the real, imaginary, and complex conjugate operators are applied element-wise.

The powers demanded by the loads at node n, denoted $\mathbf{p}^d_n$ and $\mathbf{q}^d_n$, are not necessarily treated as variables in this linear expression. In the simplest case, where all loads on the network are modeled as constant-P, they are 3x1 vectors of constants. However, in the

original LUPFM, they are defined as combination constant-P and constant-Z loads, which are expressed element-wise as:

$$\begin{aligned} p_n^{d\phi} &= p_n^{d\phi\ nom}(c_{Pn} + c_{Zn}E_n^{\phi}) \\ q_n^{d\phi} &= q_n^{d\phi\ nom}(c_{Pn} + c_{Zn}E_n^{\phi}) + C_n^{\phi} \end{aligned} \tag{2.2}$$

where $p_n^{d\phi\ nom}$, $q_n^{d\phi\ nom}$ are the active and reactive power demanded by the loads connected to node n on phase ϕ at nominal voltage and C_n is the reactive power contribution of any capacitor banks connected to node n.

Relation between voltage magnitudes and power flows

The effect of Ohm's law on voltage magnitude is enforced across each conductor in the network with Equation (2.3). First, we define quantities $\mathbf{H}_{nm}$ and $\mathbf{\Gamma}_n^{est}$.

$\forall n, m \in \mathcal{N}; nm \in \mathcal{E}$, Let

$$\mathbf{H}_{nm} = (\mathbf{Z}_{nm}\mathbf{I}_{nm}^{est}) \circ (\mathbf{Z}_{nm}\mathbf{I}_{nm}^{est})^*$$

$$\mathbf{\Gamma}_n^{est} = \begin{bmatrix} 1 & \gamma_n^{ab} & \gamma_n^{ac} \\ \gamma_n^{ba} & 1 & \gamma_n^{bc} \\ \gamma_n^{ca} & \gamma_n^{cb} & 1 \end{bmatrix}$$

These definitions come from the matrix analogs to coefficients of the binomial expansion $(\mathbf{V}_m + \mathbf{Z}_{nm}\mathbf{I}_{nm}^{est}) \circ (\mathbf{V}_m + \mathbf{Z}_{nm}\mathbf{I}_{nm}^{est})^*$. We define $\gamma_n^{\phi\psi} = V_n^{\phi est}/V_n^{\psi est}$ as the complex-valued ratio of voltages on the phases ϕ and ψ of node n. Every element of the Γ_n^{est} matrix is a constant derived from estimated values of the nodal voltage phasors, though in the expression above the *est* superscript is not included on the γ elements of the right-hand side for clarity of representation.

With these definitions, the relation between the three-phase voltage magnitude at nodes n,m and power flows through each line nm can be expressed:

$$\mathbf{E}_n = \mathbf{E}_m + 2\mathbf{M}_{nm}\mathbf{P}_{nm} - 2\mathbf{N}_{nm}\mathbf{Q}_{nm} + \mathbf{H}_{nm} \tag{2.3}$$

where

$$\mathbf{M}_{nm} = \mathbf{Re}(\mathbf{\Gamma}_m^{est} \circ \mathbf{Z}_{nm}^*)$$
$$\mathbf{N}_{nm} = \mathbf{Im}(\mathbf{\Gamma}_m^{est} \circ \mathbf{Z}_{nm}^*)$$

Relation between voltage angles and power flows

Equation (2.4) provides a second expression of Ohm's law in terms of its effects on voltage phasor angle.

$\forall n, m \in \mathcal{N}; nm \in \mathcal{E}$:

$$|\mathbf{V}_m^{est}| \circ |\mathbf{V}_n^{est}| \circ (\mathbf{\Theta}_n - \mathbf{\Theta}_m) = -\mathbf{N}_{nm}\mathbf{P}_{nm} - \mathbf{M}_{nm}\mathbf{Q}_{nm} \tag{2.4}$$

where $\mathbf{M}_{nm}$ and $\mathbf{N}_{nm}$ are defined as above.

Other Constraints

Slack bus enforcement

The "slack" or "swing" bus of a power system defines the fixed voltage magnitude and angular reference for the network when solving for the power flow. Where a substation or large centralized generator exists on the network, it becomes the obvious choice of slack bus for the system. In simulation, its voltage values are generally set to a value of 1 per unit and balanced across each phase. In the original formulation of the LUPFM, the fixed voltage at the slack bus is enforced as constraints given by Equations (2.5).

$\mathbf{E}_{slack} \in \mathcal{N}$:

$$\begin{aligned} \mathbf{E}_{slack} &= \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}^T \\ \mathbf{\Theta}_{slack} &= \begin{bmatrix} 0 & -\frac{2\pi}{3} & \frac{2\pi}{3} \end{bmatrix}^T \end{aligned} \tag{2.5}$$

Voltage magnitude constraints

In all power system applications, grid operators define upper and lower voltage bounds that nodal voltages are required to respect. These can be enforced as part of an OPF formulation with the constraints of Equation (2.6).

$\forall n \in \mathcal{N}$:

$$\begin{bmatrix} \underline{\mathrm{V}}_{RMS}^2 & \underline{\mathrm{V}}_{RMS}^2 & \underline{\mathrm{V}}_{RMS}^2 \end{bmatrix}^T \leq \mathbf{E}_n \leq \begin{bmatrix} \overline{V}_{RMS}^2 & \overline{V}_{RMS}^2 & \overline{V}_{RMS}^2 \end{bmatrix}^T \tag{2.6}$$

where $\underline{\mathrm{V}}_{RMS}$ and $\overline{V}_{RMS}$ are the lower and upper limits of the allowable RMS voltage magnitudes, respectively.

DER actuation constraints

The active and reactive power that can be dispatched by a DER is constrained by an apparent power limit. For a DER capable of four-quadrant operation on phase ϕ, this limit is circular in the plane defined by $p^{g\phi}$, $q^{g\phi}$. However, that constraint can be transformed into an arbitrary number K of linear constraints on active and reactive power as in Equation (2.7).

$\forall n \in \{s_n^\phi > 0,\ \phi \in \mathcal{P}_n\}$

$$\mathbf{\Psi}_{\cos} \circ \mathbf{p}_n^g + \mathbf{\Psi}_{\sin} \circ \mathbf{q}_n^g \leq \mathbf{s}_n^{Rated} \tag{2.7}$$

where

$$\begin{aligned} \mathbf{\Psi}_{\cos} &= \begin{bmatrix} \cos(\psi) & \cos(\psi) & \cos(\psi) \end{bmatrix}^T \\ \mathbf{\Psi}_{\sin} &= \begin{bmatrix} \sin(\psi) & \sin(\psi) & \sin(\psi) \end{bmatrix}^T \\ \psi &= \frac{2\pi k}{K} \ \ for \ \ k = 0, 1, ..., K-1. \end{aligned}$$

The generation capacity error introduced by this approximation can be made arbitrarily small at the expense of added constraints and increased computation time. When drawn in the space of active vs. reactive power, the difference in area between the set of linearized constraints and the original circular restriction on apparent power is less than 1% for $K \approx 20$ and 0.1% for $K \approx 60$ [62]. This half-space approximation is commonly used in OPF literature [1]

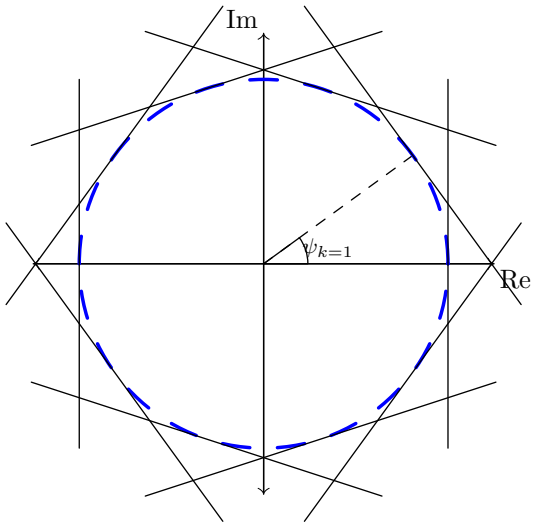

Figure 2.5: A constraint on apparent power approximated by linear half-spaces

Apparent power constraints on distribution lines were not taken into account in this model, but could be included in future work with a similar approximate treatment.

Chapter 3

Iterative Refinement for S-PBC Targets

This chapter introduces[1] a means of mitigating error associated with linearizing the power flow equations for the LUPFM of Chapter 2, which involves adapting a technique originally developed by [62]. The end goal, which will be covered in Chapter 4, will be the application of the OPF solution techniques of a centralized S-PBC to minigrids or to microgrids operating in island mode. That application, though, first requires addressing one of the primary challenges of PBC at distribution voltage levels: the need for extreme accuracy in the phasor targets generated by the S-PBC's OPF solution. As we discuss in this chapter, that accuracy can be achieved with an iterative solution method.

The iterative method relies on the successive refinement of the LUPFM with power flow solutions generated by an accompanying nonlinear Newton-Raphson (NR) solver. This requires some alterations to the LUPFM, which will be discussed in Section 3.1. Section 3.2 will introduce our chosen nonlinear solver, the Current Injection Method, and a procedure for updating its equations to treat networks that include ideal transformers. The last section, 3.3, will demonstrate the iterative method in simulation.

Unless otherwise specified, the nomenclature defined in Section 2.2 will hold throughout this chapter as well.

3.1 An Iterative Target Generation Methodology

The LUPFM discussed in Chapter 2 has proven to be an effective tool for optimizing DER power dispatch over a network. Simulated accuracy analysis shows that, even in scenarios with large power flows on a feeder, the voltage phasor values generated by a LUPFM-based OPF solution can be made to agree with those generated by a nonlinear power flow solver to within 10^{-2} per unit in magnitude and the order of 10^{-1} degrees [62]. The disagreement

[1] A partial version of this work was accepted for publication as "Iterative Linearization for Phasor-Defined Optimal Power Dispatch" in the 52nd North American Power Symposium.

between solvers decreases further when an accurate estimated system state is used to initialize the LUPFM.

This level of accuracy is more than sufficient for the majority of OPF applications at the distribution level. However, PBC is an outlier. The difficulty arises because of the extremely low impedances and small voltage angle differences between adjacent nodes of a distribution feeder under normal conditions. Given those characteristics, the specification of active and reactive power dispatch as a function of voltage phasors, which will translate into the tasking of the L-PBCs, is a highly sensitive problem.

As a motivating example: the IEEE 13-Node Test Feeder (IEEE13), which will be used as a base-case simulation throughout this chapter and the next, has a voltage base of 2.4kV line-to-ground and contains a single-phase line with an impedance of magnitude $|Z| \approx 0.1\Omega$ [35]. Relative to that base, if a set of phasor references generated by a supervisory controller contains a per-unit error of magnitude Δ in the expected difference between the voltage phasors on either side of that conductor, it will translate into the attempted recruitment of 24Δ kiloAmps of erroneous actuation from DERs throughout the network. This would likely cause those DERs to saturate in an attempt to meet their targets and result in an extremely sub-optimal state for the feeder.

It is granted that this is an extreme scenario involving a very short distribution line, but it illustrates the fact that PBC shows very little tolerance for any mismatch between the S-PBC-generated phasor targets and the closest set of network nodal voltages that is physically realizable for local controllers. In effect, for an implementation of PBC to be successful, all error associated with the approximations of the LUPFM must be eliminated.

There are many examples of algorithms that rely on iterative improvements to linearized models in order to solve OPF. Frequently, those algorithms re-initialize their linear models at each iterative step using the optimal values generated by the previous iteration, e.g. [4]. Other approaches make use of a separate, external solver, which will be the strategy employed in our method. This strategy was originally applied to the LUPFM in [62] as part of an example feeder-reconfiguration application. In that work, it was found that the iterative method significantly reduced the linearization error, but that the error in angular estimation could not be made arbitrarily small. There existed a threshold past which, it was believed, the use of the small-angle approximation in Equation (2.4) became a limiting factor in the LUPFM's accuracy.

Regardless of the specifics of the strategy used to refine the linear model, the algorithms that define most iterative linearization methods tend to follow a similar pattern, which will be described in the following subsection as it applies to the OPF problem of Chapter 2. The ultimate goal of applying iterative refinement to that problem is the driving of the LUPFM's approximated constraints to admit a set of values $\mathbf{V}_n, \mathbf{p}_n^g, \mathbf{q}_n^g \ \ \forall n \in \mathcal{N}$ that constitute a near-exact solution to the nonlinear power flow equations.

The Iterative Method Structure

In the first iteration of the method, the linearized OPF problem is solved using constraints provided by the LUPFM as specified in Equations (2.1)-(2.7), with several alterations to be discussed in the following subsection. We assume that no external sensor data is available at the process start, and so the LUPFM is initialized with the nodal voltage values balanced across phases and set to a value of 1 per unit. Line current values are initialized at 0.

The DER power dispatch, $\mathbf{p}_n^g, \mathbf{q}_n^g \quad \forall n \in \mathcal{N}$, generated by the solution of the LUPFM-based OPF is then passed to a nonlinear power flow solver. This nonlinear solver can, in theory, be of any type so long as its model agrees with the line impedances and ZIP load modeling of the LUPFM.

As part of the initialization of the nonlinear solver, $\mathbf{p}_n^g$ and $\mathbf{q}_n^g$ are incorporated as negative, constant-P loads into the ZIP-modeled demand expression of Equation (2.2), which will be expanded and redefined as Equation (3.2) in the following section. The injection defined by this combination of demand and OPF-generated power dispatch is denoted by $\mathbf{p}_n^{sp} = \mathbf{p}_n^g - \mathbf{p}_n^d$ and $\mathbf{q}_n^{sp} = \mathbf{q}_n^g - \mathbf{q}_n^d$, and it is used by the nonlinear solver to generate magnitude and phase angle values for each of the network's nodal voltages and line currents. We denote the values returned by the nonlinear solver as $\mathbf{V}_n^{NL}$ and $\mathbf{I}_{nm}^{NL}$ for all $n \in \mathcal{N}$ and $nm \in \mathcal{E}$.

At this point, the first iteration of the process is complete and the voltage phasors generated by the nonlinear solver are compared to those of the LUPFM OPF's solution. If the difference between those two sets of voltage phasors is larger than a predetermined convergence threshold, the next iteration of the process begins. Each successive iteration re-initializes the LUPFM using the nodal voltage and line current values from the nonlinear solution of the previous iteration as $\mathbf{V}_n^{est}$ and $\mathbf{I}_{nm}^{est}$ for all $n, m \in \mathcal{N}$ and $nm \in \mathcal{E}$.

This procedure is summarized below and shown graphically in Figure 3.1. Note that ε_V and ε_θ can be set to an arbitrary accuracy threshold.

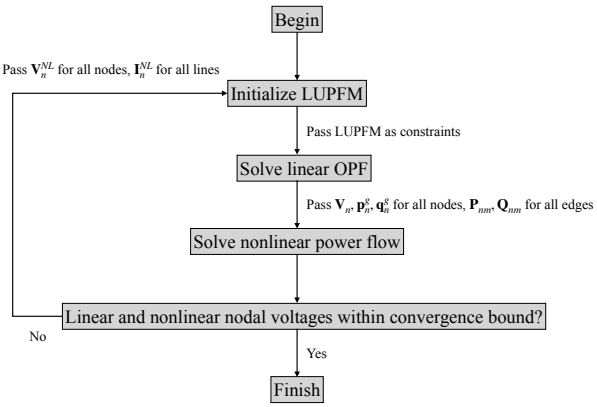

Figure 3.1: The iterative update procedure for the LUPFM OPF and the nonlinear solver

Algorithm 1: The Iterative Update Method for Refinement of Phasor Targets

for $n \in \mathcal{N}, (A, B, C) \in \mathcal{P}_n$, $nm \in \mathcal{E}, \phi \in \mathcal{P}_{nm}$ **do**

  Initialize $V_n^{A\,NL} = 1\angle 0°, V_n^{B\,NL} = 1\angle -120°, V_n^{C\,NL} = 1\angle 120°$;

  Initialize $I_{nm}^{\phi\,NL} = 0$;

end

$\mathbf{V}_n^{est}, \mathbf{I}_{nm}^{est} \leftarrow \mathbf{V}_n^{NL}, \mathbf{I}_{nm}^{NL} \quad \forall n \in \mathcal{N}, nm \in \mathcal{E}$;

Solve QP OPF described in Chapter 2 for $\mathbf{V}_n, \mathbf{p}_n^g, \mathbf{q}_n^g, \mathbf{P}_{nm}, \mathbf{Q}_{nm} \quad \forall n \in \mathcal{N}, nm \in \mathcal{E}$;

while $\max_{n\in\mathcal{N},\phi\in\mathcal{P}_n} (V_n^{\phi} - V_n^{\phi\,NL}) \leq \varepsilon_V \vee \max_{n\in\mathcal{N},\phi\in\mathcal{P}_n} (\theta_n^{\phi} - \theta_n^{\phi\,NL}) \leq \varepsilon_\theta$ **do**

  $\mathbf{p}_n^{sp}, \mathbf{q}_n^{sp} \leftarrow \mathbf{p}_n^d - \mathbf{p}_n^g, \mathbf{q}_n^d - \mathbf{q}_n^g \; \forall n \in \mathcal{N}$, where $\mathbf{p}_n^d, \mathbf{q}_n^d$ are defined as in Equation (3.2) and $\mathbf{p}_n^g, \mathbf{q}_n^g$ are treated as constant-P loads;

  Solve nonlinear power flow equations for $\mathbf{V}_n^{NL}, \mathbf{I}_{nm}^{NL}$, using load values $\mathbf{p}_n^{sp}, \mathbf{q}_n^{sp}$ and initializing with $\mathbf{V}_n, \mathbf{P}_{nm}, \mathbf{Q}_{nm} \quad \forall n \in \mathcal{N}, nm \in \mathcal{E}$;

  $\mathbf{V}_n^{est}, \mathbf{I}_{nm}^{est} \leftarrow \mathbf{V}_n^{NL}, \mathbf{I}_{nm}^{NL} \quad \forall n \in \mathcal{N}, nm \in \mathcal{E}$;

  Solve QP OPF described in Chapter 2 for $\mathbf{V}_n, \mathbf{p}_n^g, \mathbf{q}_n^g, \mathbf{P}_{nm}, \mathbf{Q}_{nm} \quad \forall n \in \mathcal{N}, nm \in \mathcal{E}$;

end

Alterations to the LUPFM

As mentioned, previous implementations of this algorithm have found that there is a lower bound beyond which the linearization error in angular estimation cannot be reduced using the original form of the LUPFM. This is corrected with a small alteration to Equation (2.4), which will be described below. Other adjustments to the LUPFM covered in this subsection include a means of introducing constant-I components of ZIP-modeled loads and a relationship that allows for tap-changing transformers to be included in the model.

Taylor Series Expansion of Angular Relationship

It is most likely the use of the small-angle approximation that keeps the expression of Equation (2.4) from being made arbitrarily accurate through iterative refinement. We replace it here with a more general first-order Taylor Series that removes the block on the ability of the OPF to approach exact solutions. That is, in all iterations after the first, Equation (2.4) is replaced by Equation (3.1).

$$|\mathbf{V}_m^{est}| \circ |\mathbf{V}_n^{est}| \circ \boldsymbol{\Delta}_{\sin\theta}^{est} + |\mathbf{V}_m^{est}| \circ |\mathbf{V}_n^{est}| \circ \boldsymbol{\Delta}_{\cos\theta}^{est} \circ [(\boldsymbol{\Theta}_n - \boldsymbol{\Theta}_m) - (\boldsymbol{\Theta}_n^{est} - \boldsymbol{\Theta}_m^{est})] = - \mathbf{N}_{nm}\mathbf{P}_{nm} - \mathbf{M}_{nm}\mathbf{Q}_{nm} \quad (3.1)$$

Where $\mathbf{\Delta}^{est}_{\cos\theta} = \begin{bmatrix}\cos(\theta^a_n - \theta^a_m)\\ \cos(\theta^b_n - \theta^b_n)\\ \cos(\theta^c_n - \theta^c_n)\end{bmatrix}$ and $\mathbf{\Delta}^{est}_{\sin\theta} = \begin{bmatrix}\sin(\theta^a_n - \theta^a_m)\\ \sin(\theta^b_n - \theta^b_n)\\ \sin(\theta^c_n - \theta^c_n)\end{bmatrix}$

All angular elements on the right-hand sides of the equations defining $\mathbf{\Delta}^{est}_{\cos\theta}$ and $\mathbf{\Delta}^{est}_{\sin\theta}$ are estimated quantities, though the superscript *est* has been dropped to avoid notational clutter.

Constant-I Load Modeling

The ZIP load model, which treats electric loads as a combination of constant-impedance, constant-current, and constant-power components, is commonly used in power systems analysis [37]. In order to faithfully model real loads using that system, it is important that all three components are capable of being represented as constraints on our OPF QP. As such, we add another Taylor Series term to Equation (2.2) that incorporates constant-I loads by expressing the voltage magnitude on phase ϕ of node n as an expansion of $\sqrt{E^{\phi}_n}$.

Equation (3.2) replaces Equation (2.2) and redefines $\mathbf{p}^{d\phi}_n$ and $\mathbf{q}^{d\phi}_n$ element-wise.

$$\begin{aligned} p^{d\phi}_n &= p^{d\phi\ nom}_n[c_P + c_I(\frac{1}{2}|V^{\phi est}_n| + \frac{1}{2|V^{\phi est}_n|}E^{\phi}_n) + c_Z E^{\phi}_n] \\ q^{d\phi}_n &= q^{d\phi\ nom}_n[c_P + c_I(\frac{1}{2}|V^{\phi est}_n| + \frac{1}{2|V^{\phi est}_n|}E^{\phi}_n) + c_Z E^{\phi}_n] \end{aligned} \tag{3.2}$$

In the above, $p^{d\phi\ nom}_n$ and $q^{d\phi\ nom}_n$ are the total active and reactive power demanded at node n when the system is operating at nominal voltage.

Transformer Modeling

There is another addition to the LUPFM that is necessary for the modeling of many feeders: the inclusion of load tap changing (LTC) transformers and voltage regulators with mechanically adjustable turns ratios. Transformers with fixed turns ratios can be accommodated in the per-unit measurement system without any special treatment beyond a consistent selection of voltage and power bases [25]. LTCs, though, require the enforcement of a voltage ratio that can be parameterized by tap setting.

Given a transformer that connects node n to a downstream node m, with tap settings that establish the 3x1 vector of turns ratios $\mathbf{r}$ between its primary and secondary sides, we can relate the primary and secondary voltages by Equation (3.3).

$$\begin{aligned} \mathbf{E}_m &= \mathbf{r}^2 \circ \mathbf{E}_n \\ \boldsymbol{\theta}_n &= \boldsymbol{\theta}_m \end{aligned} \tag{3.3}$$

where $\mathbf{r}^2 = \mathbf{r} \circ \mathbf{r}$.

The powers P_{nm} and Q_{nm} flowing through the transformer do not need to be subject to any constraints, as the power balance enforcement of Equation (2.1) will ensure that the necessary relationships are respected.

Where non-ideal transformer modeling is necessary, LTCs can be treated as a series combination of a connecting line and ideal transformer.

3.2 The Companion Nonlinear Solver

The iterative algorithm is agnostic to the specific method used to obtain the full, nonlinear power flow solution that will be used to refine the LUPFM in each step. In this work, we choose the Current Injection Method (CIM) to fill the role of the nonlinear solver.

The CIM is a popular means of solving the power flow equations on distribution networks. It was presented as an alternative to the classic NR methods of [71] and others, and it was meant to address challenges posed to those solvers by unbalanced networks with low X/R ratios. The CIM is presented in full in [24]; we summarize our reduced version, meant to treat only PQ-type buses, in this section.

Defining the CIM

Consider a network with N nodes, indexed from 1 to N. We express the nodal voltages to be solved by the CIM in Cartesian coordinates as $V_n^{\phi\,NL} = V_{n_{Re}}^{\phi} + jV_{n_{Im}}^{\phi}$. The network's nodal admittance matrix is denoted $Y_{nm}^{\phi\psi} = G_{nm}^{\phi\psi} + jB_{nm}^{\phi\psi}$, where $G_{nm}^{\phi\psi}$ and $B_{nm}^{\phi\psi}$ are the entries of the conductance and susceptance matrices for phases ϕ and ψ of nodes n and m.

The CIM begins by initializing all nodal voltage values. It then calculates the mismatch between the current demand specified at each node and the inflow or outflow calculated from the product of the admittance matrix and those nodal voltages. The real and imaginary parts of that current mismatch can be expressed as Equation (3.4).

$$\begin{aligned}
\Delta I_{n_{Re}}^{\phi} &= \frac{p_n^{sp\phi} V_{n_{Re}}^{\phi} + q_n^{sp\phi} V_{n_{Im}}^{\phi}}{(V_{n_{Re}}^{\phi})^2 + (V_{n_{Im}}^{\phi})^2} - \sum_{m\in\mathcal{N}} \sum_{\psi\in\mathcal{P}_n} (G_{nm}^{\phi\psi} V_{m_{Re}}^{\psi} - B_{nm}^{\phi\psi} V_{m_{Im}}^{\psi}) \\
\Delta I_{n_{Im}}^{\phi} &= \frac{p_n^{sp\phi} V_{n_{Im}}^{\phi} - q_n^{sp\phi} V_{n_{Re}}^{\phi}}{(V_{n_{Re}}^{\phi})^2 + (V_{n_{Im}}^{\phi})^2} - \sum_{m\in\mathcal{N}} \sum_{\psi\in\mathcal{P}_n} (G_{nm}^{\phi\psi} V_{m_{Im}}^{\psi} - B_{nm}^{\phi\psi} V_{m_{Re}}^{\psi})
\end{aligned} \tag{3.4}$$

where

$$\begin{aligned}
p_n^{sp\phi} &= p_n^{g\phi} - p_{ncP}^{d\phi} - p_{ncI}^{d\phi} |V_n^{\phi\,NL}| - p_{ncZ}^{d\phi} |V_n^{\phi\,NL}|^2 \\
q_n^{sp\phi} &= q_n^{g\phi} - q_{ncP}^{d\phi} - q_{ncI}^{d\phi} |V_n^{\phi\,NL}| - q_{ncZ}^{d\phi} |V_n^{\phi\,NL}|^2
\end{aligned}$$

are the expressions of the net active and reactive power injections specified at node n. As part of the above definition, note that we have separated the nominal demand at n into its ZIP model components:

$$\begin{aligned}
p_{ncZ}^{d\phi} + jq_{ncZ}^{d\phi} &= c_{Zn}(p_n^{d\phi\,nom} + jq_n^{d\phi\,nom}) \\
p_{ncI}^{d\phi} + jq_{ncI}^{d\phi} &= c_{In}(p_n^{d\phi\,nom} + jq_n^{d\phi\,nom}) \\
p_{ncP}^{d\phi} + jq_{ncP}^{d\phi} &= c_{Pn}(p_n^{d\phi\,nom} + jq_n^{d\phi\,nom})
\end{aligned}$$

and that the real and reactive power dispatch that were generated by the solution to the LUPFM OPF have been included in the specification as constant power loads $p_n^{g\phi}$ and $q_n^{g\phi}$

We can now define the constitutive equation of the CIM, Equation (3.5), which is to be solved using Newton's method. Let $\Delta\mathbf{I}_{n_{Re}}$ and $\Delta\mathbf{I}_{n_{Im}}$ be the 3x1 vectors expressed element-wise by Equation (3.4). $\Delta\mathbf{V}_{n_{Re}}$ and $\Delta\mathbf{V}_{n_{Im}}$ will be the adjustments made to the voltage $\mathbf{V}_n^{NL}$ at each step of the NR method. These vectors are related by

$$\Delta\mathbf{I} = \mathbf{Y}^{\dagger}\Delta\mathbf{V} \tag{3.5}$$

where

$$\mathbf{Y}^{\dagger} = \begin{bmatrix} \mathbf{B}'_{11} & \mathbf{G}'_{11} & \mathbf{B}_{12} & \mathbf{G}_{12} & \dots & \mathbf{B}_{1N} & \mathbf{G}_{1N} \\ \mathbf{G}''_{11} & -\mathbf{B}''_{11} & \mathbf{G}_{12} & -\mathbf{B}_{12} & \dots & \mathbf{G}_{1N} & -\mathbf{B}_{1N} \\ \mathbf{B}_{21} & \mathbf{G}_{21} & \mathbf{B}'_{22} & \mathbf{G}'_{22} & \dots & \mathbf{B}_{2N} & \mathbf{G}_{2N} \\ \mathbf{G}_{21} & -\mathbf{B}_{21} & \mathbf{G}''_{22} & -\mathbf{B}''_{22} & \dots & \mathbf{G}_{2N} & -\mathbf{B}_{2N} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \mathbf{B}_{N1} & \mathbf{G}_{N1} & \mathbf{B}_{N2} & \mathbf{G}_{N2} & \dots & \mathbf{B}'_{NN} & \mathbf{G}'_{NN} \\ \mathbf{G}_{N1} & -\mathbf{B}_{N1} & \mathbf{G}_{N2} & -\mathbf{B}_{N2} & \dots & \mathbf{G}''_{NN} & -\mathbf{B}''_{NN} \end{bmatrix}$$

$$\Delta\mathbf{I} = \begin{bmatrix} \Delta\mathbf{I}_{1_{Im}} \\ \Delta\mathbf{I}_{1_{Re}} \\ \Delta\mathbf{I}_{2_{Im}} \\ \Delta\mathbf{I}_{2_{Re}} \\ \vdots \\ \Delta\mathbf{I}_{N_{Im}} \\ \Delta\mathbf{I}_{N_{Re}} \end{bmatrix} \qquad \Delta\mathbf{V} = \begin{bmatrix} \Delta\mathbf{V}_{1_{Re}} \\ \Delta\mathbf{V}_{1_{Im}} \\ \Delta\mathbf{V}_{2_{Re}} \\ \Delta\mathbf{V}_{2_{Im}} \\ \vdots \\ \Delta\mathbf{V}_{N_{Re}} \\ \Delta\mathbf{V}_{N_{Im}} \end{bmatrix}$$

The off-diagonal parts $\mathbf{B}_{nm}$ and $\mathbf{G}_{nm}$ of the matrix $\mathbf{Y}^{\dagger}$ are defined identically to those of the network's admittance matrix, but the blocks on the diagonal are defined as:

$$\mathbf{B}'_{nn} = \mathbf{B}_{nn} - \begin{bmatrix} a_n^A & 0 & 0 \\ 0 & a_n^B & 0 \\ 0 & 0 & a_n^C \end{bmatrix} \quad \mathbf{G}'_{nn} = \mathbf{G}_{nn} - \begin{bmatrix} b_n^A & 0 & 0 \\ 0 & b_n^B & 0 \\ 0 & 0 & b_n^C \end{bmatrix}$$

$$\mathbf{G}''_{nn} = \mathbf{G}_{nn} - \begin{bmatrix} c_n^A & 0 & 0 \\ 0 & c_n^B & 0 \\ 0 & 0 & c_n^C \end{bmatrix} \quad \mathbf{B}''_{nn} = \mathbf{B}_{nn} - \begin{bmatrix} d_n^A & 0 & 0 \\ 0 & d_n^B & 0 \\ 0 & 0 & d_n^C \end{bmatrix}$$

where a_n^{ϕ}, b_n^{ϕ}, c_n^{ϕ}, and d_n^{ϕ} are determined by the load model at bus n:

$$a_n^{\phi} = \frac{q_{ncP}^{d\phi}[(V_{n_{Re}}^{\phi})^2 - (V_{n_{Im}}^{\phi})^2] - 2p_{ncP}^{d\phi}V_{n_{Re}}^{\phi}V_{n_{Im}}^{\phi}}{|V_n^{\phi}|^4} + \frac{p_{ncI}^{d\phi}V_{n_{Re}}^{\phi}V_{n_{Im}}^{\phi} + q_{ncI}^{d\phi}(V_{n_{Im}}^{\phi})^2}{|V_n^{\phi}|^3} + q_{ncZ}^{d\phi}$$

$$b_n^\phi = \frac{p_{ncP}^{d\phi}[(V_{n_{Re}}^\phi)^2 - (V_{n_{Im}}^\phi)^2] + 2q_{ncP}^{d\phi}V_{n_{Re}}^\phi V_{n_{Im}}^\phi}{|V_n^\phi|^4} - \frac{q_{ncI}^{d\phi}V_{n_{Re}}^\phi V_{n_{Im}}^\phi + p_{ncI}^{d\phi}(V_{n_{Re}}^\phi)^2}{|V_n^\phi|^3} - p_{ncZ}^{d\phi}$$

$$c_n^\phi = \frac{p_{ncP}^{d\phi}[(V_{n_{Im}}^\phi)^2 - (V_{n_{Re}}^\phi)^2] - 2q_{ncP}^{d\phi}V_{n_{Re}}^\phi V_{n_{Im}}^\phi}{|V_n^\phi|^4} + \frac{q_{ncI}^{d\phi}V_{n_{Re}}^\phi V_{n_{Im}}^\phi - p_{ncI}^{d\phi}(V_{n_{Im}}^\phi)^2}{|V_n^\phi|^3} - p_{ncZ}^{d\phi}$$

$$d_n^\phi = \frac{q_{ncP}^{d\phi}[(V_{n_{Re}}^\phi)^2 - (V_{n_{Im}}^\phi)^2] - 2p_{ncP}^{d\phi}V_{n_{Re}}^\phi V_{n_{Im}}^\phi}{|V_n^\phi|^4} + \frac{p_{ncI}^{d\phi}V_{n_{Re}}^\phi V_{n_{Im}}^\phi - q_{ncI}^{d\phi}(V_{n_{Re}}^\phi)^2}{|V_n^\phi|^3} - q_{ncZ}^{d\phi}$$

Adapting the CIM for LTCs

The CIM method as presented above does not take into account the voltage adjustments caused by LTCs with non-zero tap positions, and so it needs to be adjusted before it can be used to solve power flow on more general networks. This adjustment can be carried out as a set of manipulations to be included in the nonlinear solver's initialization at each step of the broader iterative method.

Let $\mathcal{T} \in \mathcal{E}$ be the set of all LTCs on the network. Let $nm \in \mathcal{T}$ be defined as the LTC connecting node n to a downstream node m. The tap-set relationship between their voltages is then given by $\mathbf{V}_m = \mathbf{r}_{nm} \circ \mathbf{V}_n$. For each member of $\mathcal{T}$ we introduce the new quantity $\mathbf{I}_{nm}^{NL} = \mathbf{I}_{nm_{Re}} + j\mathbf{I}_{nm_{Im}}$, which we define as the current entering the transformer's downstream node m. The value $\Delta\mathbf{I}_{nm_{Re}} + j\Delta\mathbf{I}_{nm_{Im}}$ will be the adjustment to that quantity made at each update step of the NR method.

We use that quantity to redefine Equation (3.4) as Equation (3.6), so that it takes into account the current entering and leaving each node through LTCs. This is a convenient means of tracking the current that flows across the ideal transformers in the network, but it is not strictly necessary that it be included as a part of the CIM. On large networks with many LTCs, it may be preferable to leave Equation (3.4) unchanged and to solve for nodal voltages alone.

$$\begin{aligned}
\Delta I_{n_{Re}}^\phi = \frac{p_n^{sp\phi}V_{n_{Re}}^\phi + q_n^{sp\phi}V_{n_{Im}}^\phi}{(V_{n_{Re}}^\phi)^2 + (V_{n_{Im}}^\phi)^2} &- \sum_{m\in\mathcal{N}, nm\notin\mathcal{T}} \sum_{\psi\in\mathcal{P}_n} (G_{nm}^{\phi\psi}V_{m_{Re}}^\psi - B_{nm}^{\phi\psi}V_{m_{Im}}^\psi) \\
&- \sum_{nm\in\mathcal{T}} \sum_{\psi\in\mathcal{P}_n} r_{nm}^\phi \mathbf{I}_{nm_{Re}} + \sum_{mn\in\mathcal{T}} \sum_{\psi\in\mathcal{P}_n} \mathbf{I}_{mn_{Re}} \\
\Delta I_{n_{Im}}^\phi = \frac{p_n^{sp\phi}V_{n_{Im}}^\phi - q_n^{sp\phi}V_{n_{Re}}^\phi}{(V_{n_{Re}}^\phi)^2 + (V_{n_{Im}}^\phi)^2} &- \sum_{m\in\mathcal{N}, nm\notin\mathcal{T}} \sum_{\psi\in\mathcal{P}_n} (G_{nm}^{\phi\psi}V_{m_{Im}}^\psi - B_{nm}^{\phi\psi}V_{m_{Re}}^\psi) \\
&- \sum_{nm\in\mathcal{T}} \sum_{\psi\in\mathcal{P}_n} r_{nm}^\phi \mathbf{I}_{nm_{Im}} + \sum_{mn\in\mathcal{T}} \sum_{\psi\in\mathcal{P}_n} \mathbf{I}_{mn_{Im}}
\end{aligned} \tag{3.6}$$

We next make adjustments to the entries of Equation (3.5). In order to do that, we first must introduce several definitions:

Let $\mathbf{Y}^{\dagger}_{col_k}$ be the submatrix of $\mathbf{Y}^{\dagger}$ made up of the columns associated with the kth node of the system, i.e.

$$\mathbf{Y}^{\dagger}_{col_k} = \begin{bmatrix} \mathbf{Y}^{\dagger}\mathbf{e}_{6(k-1)+1} & \mathbf{Y}^{\dagger}\mathbf{e}_{6(k-1)+2} & \mathbf{Y}^{\dagger}\mathbf{e}_{6(k-1)+3} & \mathbf{Y}^{\dagger}\mathbf{e}_{6(k-1)+4} & \mathbf{Y}^{\dagger}\mathbf{e}_{6(k-1)+5} & \mathbf{Y}^{\dagger}\mathbf{e}_{6(k-1)+6} \end{bmatrix}$$

where $\mathbf{e}_i$ is the ith elementary vector of the $6N$x$6N$ matrix.

Let $\mathbf{Y}^{\dagger}_{sub_{\ell,k}}$ be the 6x6 submatrix of $\mathbf{Y}^{\dagger}$ associated with the ℓth and kth nodes:

$$\mathbf{Y}^{\dagger}_{sub_{\ell,k}[i,j]} = \mathbf{Y}^{\dagger}_{[6(\ell-1)+i,6(k-1)+j]} \;\; \forall\, 1 \leq i,j \leq 6$$

Lastly, let $\mathbf{diag}(\lambda)$ represent the pxp diagonal matrix formed from the entries of an arbitrary px1 vector λ:

$$\mathbf{diag}(\lambda) = \sum_{i=1}^{p} (\mathbf{e}_i^T \lambda)\mathbf{e}_i\mathbf{e}_i^T$$

The CIM method for solving power flow can then be redefined as in the following algorithm. ε can be set to an arbitrary accuracy threshold.

Algorithm 2: The CIM Algorithm, Adapted for LTCs

Initialize $\mathbf{V}_n^{NL} = \mathbf{V}_n^{LUPFM} \quad \forall n \in \mathcal{N}$, where $\mathbf{V}_n^{LUPFM}$ are the results of the most recent solution of the OPF QP;

Initialize $I_{nm}^{\phi\, NL} = (P_{nm}^{\phi\, LUPFM} + jQ_{nm}^{\phi\, LUPFM})/V_n^{\phi\, LUPFM} \quad \forall nm \in \mathcal{T}, \phi \in \mathcal{P}_{nm}$;

Initialize $\Delta\mathbf{V}_{n_{Re}} = \infty$, $\Delta\mathbf{V}_{n_{Im}} = \infty$;

while $\max(\Delta\mathbf{V}_{n_{Re}}) \geq \varepsilon \vee \max(\Delta\mathbf{V}_{n_{Im}}) \geq \varepsilon$ **do**

  Solve Equation (3.6) for $\Delta\mathbf{I}_{n_{Re}}$, $\Delta\mathbf{I}_{n_{Im}}$; Build $\mathbf{Y}^{\dagger}$ as in Equation (3.5);

  for $nm \in \mathcal{T}, \phi \in \mathcal{P}_{nm}$ **do**

    Replace $\mathbf{Y}^{\dagger}_{col_n}$ with $\mathbf{Y}^{\dagger}_{col_n} + \mathbf{Y}^{\dagger}_{col_m} \cdot \mathbf{diag}(\begin{bmatrix}\mathbf{r}_{nm}\\ \mathbf{r}_{nm}\end{bmatrix})$;

    Replace all entries of $\mathbf{Y}^{\dagger}_{col_m}$ with 0;

    $\mathbf{Y}^{\dagger}_{sub_{m,m}} \leftarrow -\mathbf{I}_{6x6}$, the negative 6x6 identity matrix;

    $\mathbf{Y}^{\dagger}_{sub_{n,m}} \leftarrow \mathbf{diag}(\begin{bmatrix}\mathbf{r}_{nm}\\ \mathbf{r}_{nm}\end{bmatrix})$;

    Redefine the $\Delta\mathbf{V}$ vector of Equation (3.5) by replacing $\Delta\mathbf{V}_{m_{Re}}$ and $\Delta\mathbf{V}_{m_{Im}}$ with $\Delta\mathbf{I}_{nm_{Im}}$ and $\Delta\mathbf{I}_{nm_{Re}}$, respectively;

  end

  Solve Equation (3.5) using the adjusted versions of $\mathbf{Y}^{\dagger}$ and $\Delta\mathbf{V}$;

  $\mathbf{V}_n^{NL} \leftarrow \mathbf{V}_n^{NL} + \Delta\mathbf{V}_{n_{Re}} + j\Delta\mathbf{V}_{n_{Im}} \quad \forall n \in \mathcal{N}$;

  $\mathbf{I}_{nm}^{NL} \leftarrow \mathbf{I}_{nm}^{NL} + \Delta\mathbf{I}_{nm_{Re}} + j\Delta\mathbf{I}_{nm_{Im}} \forall nm \in \mathcal{T}$;

  $\mathbf{V}_m^{NL} \leftarrow \mathbf{r}_{nm} \cdot \mathbf{V}_n^{NL} \quad \forall nm \in \mathcal{T}$;

end

Essentially, this algorithm removes the voltages of all nodes to which a transformer's secondary side is connected from Equation (3.5), as those values are now completely specified by tap relation to the voltage on the transformer's primary side.

The corresponding elements $\mathbf{\Delta V}$ are then redefined as updates to the transformer's current, $\Delta\mathbf{I}_{nm_{Re}}$ and $\Delta\mathbf{I}_{nm_{Im}}$ for $nm \in \mathcal{T}$. As mentioned earlier, computation-time concerns may make it more desirable to solve for voltages alone with the CIM and to calculate transformer currents after the fact. If that is the case, Equation (3.6) should be replaced with Equation (3.4) in Algorithm 2 and the steps associated with $\mathbf{Y}^{\dagger}_{sub_{m,m}}$,$\mathbf{Y}^{\dagger}_{sub_{n,m}}$, and $\mathbf{I}_{nm}^{NL}$ should be ignored. The variables $\mathbf{\Delta V}_m, \mathbf{\Delta I}_m$ associated with the transformer secondary sides would then be removed without replacement in lieu of the redefinition of $\mathbf{\Delta V}$ and any loads connected directly to the secondary sides of transformers would be referenced to the primary sides.

The relationship between transformer primary and secondary voltages is reflected in $\mathbf{Y}^{\dagger}$ with the multiplication of the columns associated with each transformer's secondary node

m by a diagonal matrix of tap settings, followed by their addition to the columns associated with the transformer's primary node n. For any transformers across which currents are to be calculated, the columns of $\mathbf{Y}^{\dagger}$ associated with m are replaced with zeros.

To complete the current calculation, the 6x6 submatrix on the block diagonal of $\mathbf{Y}^{\dagger}$ associated with m is replaced with the negative identity matrix and the off-diagonal block submatrix relating n to m is replaced with the diagonal matrix of tap settings. This tracks the current entering node m and the current injection at node n, respectively. With these changes made, the resulting Equation (3.5) accurately represents the power flow relationships on the network and can be solved with standard NR methods.

3.3 Demonstrating the Iterative Method

The Simulation Enviroment

The iterative method described in Section 3.1 was tested in simulation on the IEEE 13-Node Test Feeder (IEEE13). This small, 4.16kV feeder has relatively large and unbalanced loading; it is frequently used in tests of distribution analysis strategies. It has a tap-changing transformer near its assumed point of connection to the transmission grid, and a mixture of one-, two-, and three-phase lines. The complete specification of the feeder can be found in [35].

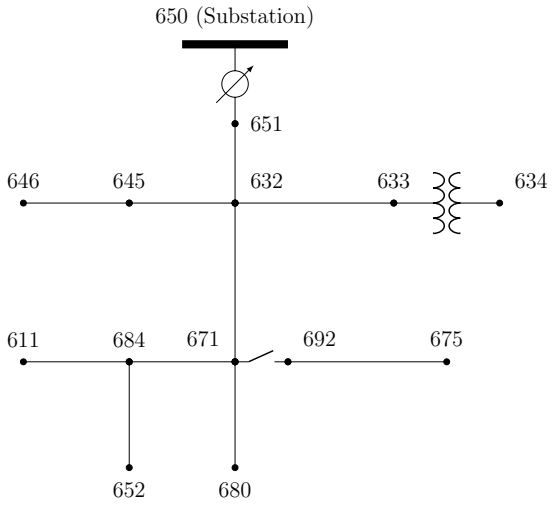

Figure 3.2: The IEEE 13-Node Feeder, one-line diagram

From its original form, we made several simplifications designed to minimize the portion of the programming implementation that is not necessary for the demonstration of core ideas:

- The voltage regulator between Bus 632 and the point of connection to the transmission grid at Bus 650 is considered to be ideal. A new node, labeled "Bus 651," is added to the secondary side of the regulator.
- The transformer XFM-1, connecting Bus 633 and Bus 634, is also considered to be ideal.
- The switch connecting Bus 671 and Bus 692 was replaced with a line segment of configuration 601 and length 0.01mi
- All delta-connected spot loads were replaced with wye-connected loads of the same magnitude. The ZIP model coefficients for each bus were kept consistent with [35].
- The distributed load on the line connecting Bus 632 and Bus 671 was replaced with a spot load of the same total value at Bus 632.
- All capacitor banks were removed from the system.

The simulations described in this section were carried out in Python. The cvxpy library's default Operator Splitting QP Solver was used for the OPF stage of each iteration. For the nonlinear solver, we have tested both OpenDSS, via the OpenDSSdirect Python library, and our own implementation of the adapted CIM method covered in Section 3.2. The following results were generated using the CIM solver.

The two operational objectives that are used as simulation cases are those introduced in Chapter 2: three-phase balancing and phasor matching. However, the notional motivation for the phasor matching objective is different than previously discussed, as we assume in this section that the IEEE13 feeder is not islanded. In this test case, rather than attempting to equalize voltages across a switch before reconnecting to a transmission grid, the phasor matching case might be used to reconfigure a feeder as described in [62].

The constraints on the OPF QP were defined as specified by the LUPFM of Equations (2.1) through (2.7), along with (3.3). Ten half-spaces were used to form the approximation of DER power limits given by Equation (2.7), and the voltage bounds of Equation (2.6) were set to 0.9 and 1.1 per unit. That is a wider allowable voltage range than would be seen in normal distribution operations; it reflects the relatively low DER penetration on the feeder and the lack of capacitor banks.

We adapt the LUPFM as described in Section 3.1: Equation (2.4) is replaced by (3.1) and (2.2) is replaced by (3.2). Adjustments to the CIM to allow for the incorporation of LTCs were made as in Section 3.2.

All constraints were identical for both the phasor matching and three-phase balancing simulations; only the objective function was changed between cases.

Simulation I: Three-phase Balancing

The three-phase balancing objective is designed to recruit DERs in an attempt to balance voltages across phases A, B, and C at each node of the feeder. The IEEE13 feeder has

heavily unbalanced loading, which is reflected in a voltage profile that is unbalanced even with the feeder's connection to the transmission network and the mitigating actions of LTC tap changes. We generated a S-PBC phasor profile designed to ensure that DERs on the feeder attempt to serve load locally and drive nodal voltages toward balanced conditions that could allow for better operation of three-phase loads.

For this test case, we added the following generators to the network:

- DERs with capacity values of 0.015 per unit, co-located with each connected load.

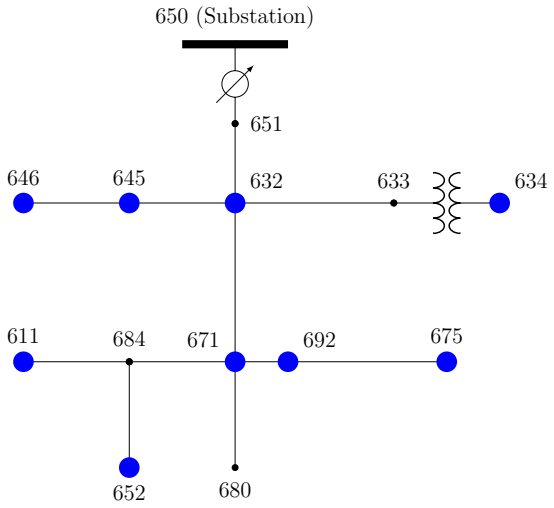

Figure 3.3: DER distribution for the three-phase balancing test

In defining our objective, we consider "perfectly balanced" conditions at a two- or three-phase node to be the state in which all voltage phasors have equal magnitude and there is 120° of angular separation between all phases $\phi \in \mathcal{P}_n$. In this definition, balanced feeder operation does not imply any restrictions on single-phase nodes.

The objective is formally defined by Equation (3.7).

$\forall n \in \mathcal{N}; \forall nm \in \mathcal{E}; \forall \phi, \psi \in \mathcal{P}_n$:

$$\min_{\mathbf{E}_n, \boldsymbol{\theta}_n, \mathbf{p}^g_n, \mathbf{q}^g_n, \mathbf{P}_{nm}, \mathbf{Q}_{nm}} \quad \sum_{\forall n \in \mathcal{P}^3 \cup \mathcal{P}^2, \phi \neq \psi} (E^\phi_n - E^\psi_n)^2 + \sum_{\forall n \in \mathcal{P}^3 \cup \mathcal{P}^2, \phi \neq \psi} (\theta^\phi_n - \theta^\psi_n \pm \frac{2\pi}{3})^2 \tag{3.7}$$

where the final form of the $\pm$ operator is determined by the expected relationships of the phases A, B, and C:

$$\theta^A_n - \theta^C_n = \theta^C_n - \theta^B_n = \theta^B_n - \theta^A_n = -\frac{2\pi}{3}$$

OPF Objective II: Phasor Matching

In the phasor-matching use case, DERs are dispatched so as to drive the voltage at a selected node $\mathbf{V}_n$ to an externally provided phasor value, $\mathbf{V}_{ext}$, which could be the voltage on the external side of a switch to be closed during grid reconfiguration. The objective is defined so as to to generate a set of voltage phasor targets that includes a value $\mathbf{V}_n$ that matches $\mathbf{V}_{ext}$.

For this test case, we added the following generators to the network:

- DERs with capacity values of 0.015 per unit. These DERs were co-located with the loads connected at all nodes excepting 671, which was the node selected for phasor matching.
- A large DER, with capacity 0.2 per unit, on all three phases of target node 671. This is meant to ensure that available generation is sufficient to achieve operational objectives in the absence of approximation error.

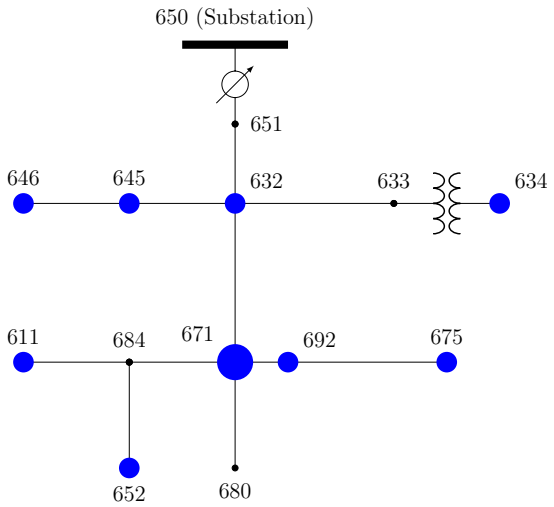

Figure 3.4: DER distribution for the phasor matching test

The phasor matching objective is defined by Equation (3.8):

$\forall n \in \mathcal{N}; \forall nm \in \mathcal{E}; \forall \phi, \psi \in \mathcal{P}_n$:

$$\min_{\mathbf{E}_n, \mathbf{\Theta}_n, \mathbf{p}^g_n, \mathbf{q}^g_n, \mathbf{P}_{nm}, \mathbf{Q}_{nm}} \quad (\mathbf{E}_{671} - \mathbf{E}_{ext})^2 + (\mathbf{\Theta}_{671} - \mathbf{\Theta}_{ext})^2 \tag{3.8}$$

Results

Three-Phase Balancing Results

Figure 3.5 plots the maximum discrepancy between the LUPFM and CIM phasors across the network over ten iterations of our solution method. The mismatch value for the first iteration is $4.93 \cdot 10^{-3}$ per unit magnitude and 0.19° of angle. Even these initial errors, generated by a LUPFM initialized to a flat start, are small. However, when translated to the real and reactive power injections demanded of DERs under the control of L-PBCs, they would result in significant discrepancies.

Over subsequent iterations the error decreases sharply, becoming $1.15 \cdot 10^{-11}$ per unit magnitude and $2.21 \cdot 10^{-10}$ degrees of angle after the tenth iteration. At that point, the linearization error is within the acceptable range of PBC, as well as any other reasonable OPF use case.

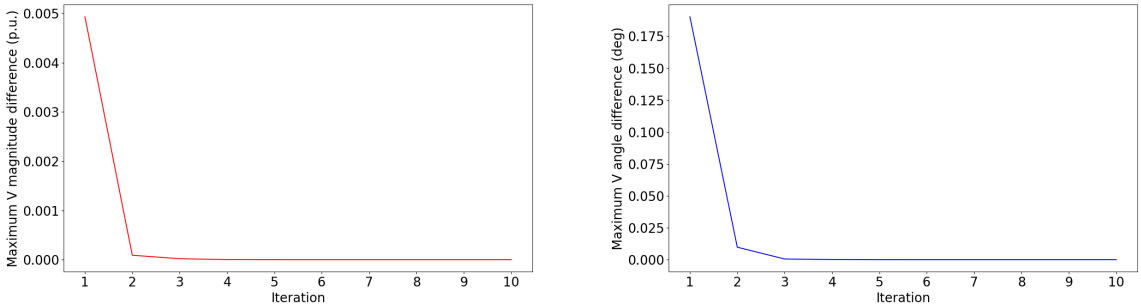

Figure 3.5: Three-phase balancing objective, maximum mismatch in magnitude (left) and angle (right) between LUPFM OPF and CIM phasor values

Figure 3.6 plots the agreement between the value of the objective function minimized by the LUPFM OPF and the value of that same objective calculated from the phasors generated by the CIM. The difference between the two decreases from $1.22 \cdot 10^{-3}$ for the first iteration to $1.74 \cdot 10^{-12}$ for the tenth.

As a metric for the phasor profile's performance with regards to its goal, we use the level of voltage balance across all three-phase nodes on the feeder. We calculate the imbalance at a node n as the magnitude ratio of the negative- and positive-sequence components[2] of the voltage: $|V_{2n}|/|V_{1n}|$. This is a definition used by the IEC [60].

With this metric, the average imbalance across all three-phase nodes is 0.39%, with a maximum of 0.62%. This can be compared to the average imbalance in the absence of PBC

[2]Positive-, negative-, and zero-sequence, or "symmetrical" components more generally, are an analytical tool introduced by Charles Fortescue in the early 1900s [22]. The concept is tangential to the material of this dissertation and will not be defined here, but readers looking for more information are encouraged to refer to [25].

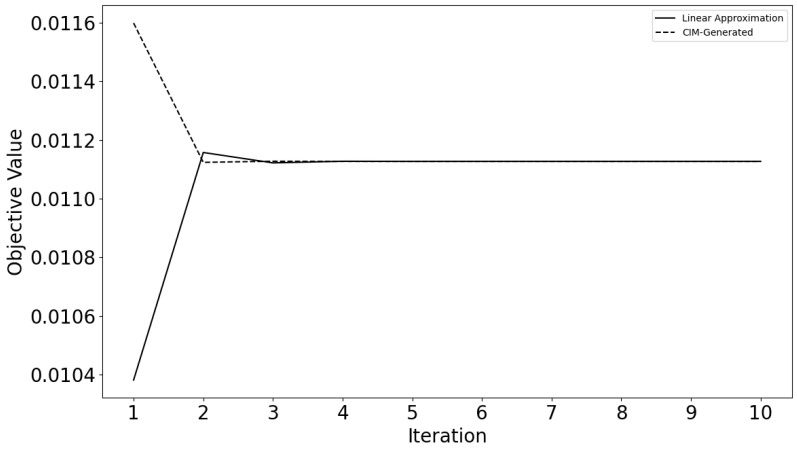

Figure 3.6: Convergence of the three-phase balancing objective values achieved by the linear and nonlinear voltage profiles

on our feeder, 0.85%. The maximum imbalance in the uncontrolled case, 1.58%, enters the range at which it is recommended that induction motors be derated.

Phasor Matching Results

As mentioned, we select Bus 671 as our target node.

$$\mathbf{V}_{ext} = \begin{bmatrix} 0.975\angle 0^\circ & 0.975\angle -120^\circ & 0.975\angle 120^\circ \end{bmatrix}^T$$

is chosen to be the phasor target value to be matched.

The results of this simulation are collected in the tables below. Table 3.1 compares the values of the OPF- and CIM-generated voltage phasors at the target bus after the first iteration of the method. Table 3.2 does the same after the tenth iteration.

	Phase A	***Phase B***	***Phase C***
V_{RMS}^{LUPFM}	0.975	0.975	0.975
θ^{LUPFM}	$1.02752527 \cdot 10^{-9}$	-120	120
V_{RMS}^{CIM}	0.97066777	0.97202337	0.97038812
θ^{CIM}	$-7.69399178 \cdot 10^{-3}$	-119.965065	120.028590

Table 3.1: Voltage Phasors at Bus 671, First Iteration

	Phase A	***Phase B***	***Phase C***
V_{RMS}^{LUPFM}	0.975	0.975	0.975
θ^{LUPFM}	$3.90110312 \cdot 10^{-9}$	-120	120
V_{RMS}^{CIM}	0.975	0.975	0.975
θ_{RMS}^{CIM}	$9.28970765 \cdot 10^{-10}$	-120	120

Table 3.2: Voltage Phasors at Bus 671, Tenth Iteration

We see a similar error reduction at the system nodes whose voltage values were not included in the objective function. Figure 3.7 shows that decrease over ten iterations of the solution process. As in Figure 3.5, we plot the maximum of the magnitude and angle differences between cvxpy- and CIM-generated phasors at all nodes in the system for each iterative step.

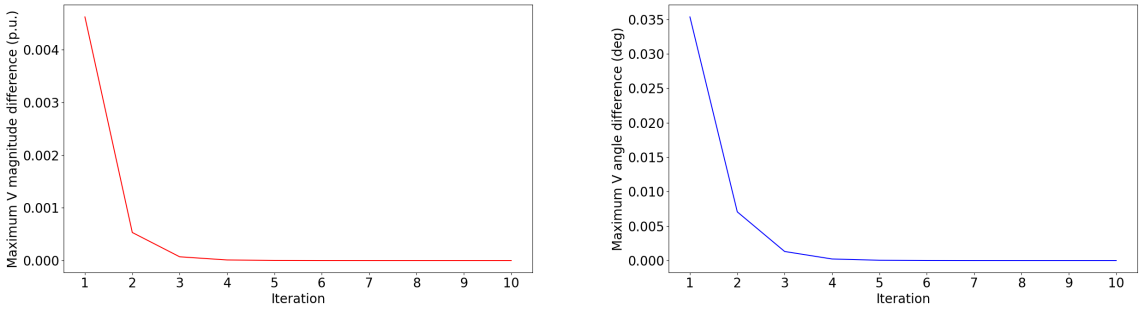

Figure 3.7: Phasor matching objective, maximum mismatch in magnitude (left) and angle (right) between LUPFM OPF and CIM phasor values

Ten refinements to the LUPFM decrease the mismatch between the phasor values to the order of 10^{-10} per unit in magnitude and 10^{-8} degrees in angle.

As in the case of three-phase balancing, the voltage phasor profile generated by this test is accurate enough to be considered feasible on a nonlinear power system. As such, it forms a suitable set of voltage phasor targets that can be communicated to L-PBCs on a distribution network.

Chapter 4

Extension of PBC to Islanded Networks

4.1 Island-Mode Considerations

The iterative method of Section 3.1 performs well on distribution systems, reliably arriving at a phasor profile that can be confirmed feasible by a nonlinear solver. When translating this implementation of PBC to islanded microgrids, though, there are additional considerations beyond linearization error that must be taken into account.

Islanded networks pose no real challenge to LUPFM-based OPF. The only necessary adjustment is the removal of Equation (2.5), the slack bus enforcement expression, from the model's constraints. Where complications arise is in finding the nonlinear power flow solution to be used in updating the LUPFM at each iteration of our solution method.

In particular, a nonlinear solver on an islanded network no longer has an obvious choice of slack bus. Because the appropriate choice of slack bus can influence the optimality of the solution, and even the ability of a numerical method to converge to a solution at all, it is necessary to have a dependable strategy for its selection. After a brief definition of the slack bus and its role in NR power flow solution algorithms, we discuss such a strategy and define an end-to-end implementation of the S-PBC. The effectiveness of this implementation is then demonstrated over a variety of simulated DER distributions.

The Slack Bus in Newton-Raphson Power Flow

The determination of a "power flow solution" was the end goal of the CIM solver introduced in Section 3.2. To formalize that term's definition, consider a set of electric loads with complex-valued demand, a network model, and some number of generating resources. Finding a power flow solution to that arrangement refers to the determination of a set of complex-valued, steady-state nodal voltage phasors that allow for all connected loads to receive the power they demand while respecting all constraints on generators. Essentially, a

power flow solution represents a feasible steady-state AC operating point for an electrical network.

Solving the equations that govern power flow requires a numerical method. NR-based techniques are frequently chosen for the application, though branch-based methods such as the backward-forward sweep (BFS) form another commonly used set of approaches, e.g. [48]. For the family of NR power flow solutions, the specific equations used with the method can vary; the classic formulation is derived from the expression for complex power [25], though the CIM used in this work [24] has also become popular for use in distribution analysis.

NR power flow methods generally divide the nodes of a system into three categories: PQ buses, at which the active and reactive power demand is considered a fixed quantity and the complex voltage must be determined; PV buses, at which the active power and the magnitude of the complex voltage are specified and the reactive power and voltage angle must be determined; and the slack bus, at which the complex voltage is fixed and the real and reactive powers must be determined. Nodes at which DERs are connected can be represented as either PV or PQ buses with negative values of power demand.

A network is allowed to have any number of PQ and PV buses, but generally only a single slack bus. The role of that slack bus is the establishment of a reference voltage angle for the system, as well as the provision of flexibility in the network-wide active and reactive power production. This flexibility is necessary to allow for an accounting of the fact that the pre-specified power injections at the fixed PQ and PV buses do not give a complete picture of the power transfer between nodes. Losses in the conductors also factor into the balance of supply and demand, but those are not known until the voltage profile across the network has been established. It is the update of the slack bus's estimated power injection that acts to cover the changes in line losses caused by updates to the network's voltage profile.

When solving for power flow on distribution systems under normal conditions, the selection of the slack bus is straightforward: it is the point at which the network connects to the larger transmission grid. Because the distribution system has very little ability to affect the voltage at its point of connection to the transmission level, that point can be considered a stiff voltage source. We can assume, as well, that active and reactive power generation constraints at that bus are not a concern, as the contribution of any given distribution feeder to the total demand on a transmission system is negligible. For the purposes of the NR solver, this means that the transmission system is capable of sourcing any amount of power necessary to account for mismatches between local feeder load and DER generation.

Choosing a Slack Bus in Island-Mode Operation

In the absence of a connection to the transmission level, no stiff voltage bus exists on a distribution network. Solutions to the power flow equations on islanded networks, then, require that the classical NR techniques be adapted. A similar problem exists for branch-based power flow methods, which generally also require a slack bus.

A number of approaches to the problem of slack bus definition for both NR and branch-based methods on islanded networks have been proposed over the past decades. Most

straightforwardly, a DER bus can be assigned to fill the slack bus role, as was done in [34]. Additional DER modeling can be used to update power dispatch values at non-slack generator buses as well [50]. This has been demonstrated effectively in simulation, but does not address cases in which there is no bus with a large enough DER that it can be confidently considered a stiff voltage source.

Another family of approaches attempts to allocate losses across multiple generating nodes. This is generally described as the "distributed" slack bus approach, and it can be traced back to the 1970s. The distributed slack bus definition involves assigning a participation factor to each generator on a network and then using it to apportion a measurement related to system losses. Various definitions of the participation factor and methods for calculating the system loss metric have been explored in a number of papers, of which [15] has an excellent review.

Many of those works that explore the distributed slack bus concept are primarily concerned with the transmission network, but efforts have been made to apply it to distribution networks as well. One notable example is [65], in which generator domains and power flow directions are analyzed in an attempt to assign loads on particular branches to the DERs most relevant to serving those portions of the network. This localized exchange is then incorporated into an iterative update scheme for participation factors that determine the proportion of feeder losses that each DER is to absorb.

While the distributed slack bus technique is interesting, the correct apportionment of losses is less critical to the nonlinear solver used in our method because of its iterative exchange with the LUPFM-based OPF. At each iteration of the process after the first, the OPF's dispatch will begin taking feeder losses into account and adjusting all controllable DERs to optimally distribute those losses across them.

A third category of approach to solving power flow on islanded networks involves introducing new physical variables, such as system frequency, that are calculated alongside the bus voltages. These new variables are used to adjust DER values according a control relationship. Examples of this include [47] for the NR method and [18] for the BFS, both of which remove the need for a classical slack bus by introducing droop-controlled DER models to the solution method. A similar approach is used in [2], which also makes use of a Newton Trust Region algorithm to solve the power flow equations, noting the convergence difficulties of NR-based methods due to the small region of attraction of islanded microgrids. This concern, too, is less pressing for our CIM solver because of the warm start initialization that it receives from the LUPFM OPF.

In our implementation of the CIM, we avoid any such explicit modeling of DER-level control schemes. This is necessary to ensure the separation between S-PBC and L-PBC layers that is called for under the PBC model; recall that the L-PBC is required to have the freedom to track its phasor target using any type of control strategy. However, we also do not assume that there exists a single bus at which sufficient DER is installed across all phases to make it an appropriate slack bus selection. We reconcile these two decisions in the CIM solver by splitting the slack bus, assigning it on a per-phase basis and potentially designating different slack buses for each of phase A, B, and C.

In making the slack bus selection on each phase, there are two primary concerns that need to be addressed:

- Sufficient generation capacity should be present to account for power losses on the relevant phase.
- The slack bus should be located as close as possible to the loads on the relevant phase.

To satisfy the first of these criteria, we calculate the sum of the network-wide estimated line losses from the voltage phasor profile generated by the LUPFM OPF, and denote that value $s_{loss}^{\phi\, LUPFM}$. That quantity is then used to define a minimum threshold for available generation capacity that would ideally be met by the DER connected to any prospective slack bus.

We implement this threshold check as a simple measure of leftover capacity in the DER connected to a bus, as determined by the solution to the LUPFM OPF. This value is then compared to the difference between $s_{loss}^{\phi\, LUPFM}$ and any losses that were already accounted for in the LUPFM's initialization during the present iteration. Formally, we describe this comparison by stating that any bus n with DER satisfying the condition

$$|s_n^{R\phi} - (p_n^{g\phi} + jq_n^{g\phi})| \geq |s_{loss}^{\phi\, LUPFM} - (\sum_{\forall km \in \mathcal{E}} P_{Lkm}^{est\phi} + jQ_{Lkm}^{est\phi})| \tag{4.1}$$

is considered to be a slack bus candidate for phase ϕ.

We note that any linearization error that generates an artificially low choice of $s_{loss}^{\phi\, LUPFM}$ could result in the selection of a slack bus without enough DER capacity to serve the losses on the feeder. However, this is not especially concerning for our iterative solution method. Because DER generation constraints on the slack bus are not enforced during the operation of the CIM, the ability of the CIM to converge to a solution will not be affected. If the CIM's solution then violates DER generation constraints at the slack bus, those constraints will be re-asserted during the next iteration of the LUPFM OPF.

Given this repeated enforcement on the part of the OPF, we can have confidence that the generation limits will ultimately be respected by our end solution. As such, the sufficient-capacity criterion can be waived if no DER can provide a satisfactory loss margin due to their previous dispatch commitments. If no bus meets the condition of Equation (4.1), that criterion is abandoned and all buses with attached DERs are evaluated as slack bus candidates.

From the set of candidate slack buses we select the node of minimal weighted electrical distance, in terms of impedance magnitude, to the feeder's loads. We define a suitability score, χ_n^{ϕ}, for all buses n on phase ϕ with Equation (4.2).

$$\chi_n^{\phi} = \sum_{\forall m \neq n \in \mathcal{N}} |Z_{nm}^{\phi\, eff}||p_m^{d\phi} + jq_m^{d\phi}| \tag{4.2}$$

where $Z_{nm}^{\phi\,eff}$ is the effective single-phase impedance between nodes n and m on phase ϕ. The bus with the lowest χ_n^ϕ score among the candidates that were selected by the sufficient-DER filter is then chosen as the slack bus for phase ϕ.

Once the slack bus has been assigned, it will be initialized at the start of the CIM to the value of the voltage phasor V_n^ϕ generated by the LUPFM OPF. This assignment and re-initialization happens at each iteration of the broader solution method, and so it is possible for the slack bus to be re-assigned between iterations. In practice, this often happens once, between the first and second iterations of the method. This is due to the fact that, after the first iteration, $\mathbf{P}_{Lkm}^{est}$ and $\mathbf{Q}_{Lkm}^{est}$ take on their first nonzero values and alter the sufficient-capacity criterion of Equation (4.1). Any further adjustments to $\mathbf{P}_{Lkm}^{est}$ and $\mathbf{Q}_{Lkm}^{est}$ introduced in successive iterations after the first are relatively minor, and they rarely alter the slack bus placement.

4.2 Demonstrating Island-Mode Operation

The simulation results presented in this section represent the culmination of this work: the demonstration of a means of applying the S-PBC component of a PBC scheme to a medium-voltage minigrid or a microgrid operating in island mode.

Simulation Setup

We again use the IEEE13 feeder, modified as in Section 3.3, for our testing, though for the simulations of this section we remove the notional connection of Bus 650 to a substation. In addition, we set the taps on all phases of the LTC connecting Bus 650 to Bus 651 to the zero position. Using that topology, we explore our controller's performance on a number of randomized DER distributions.

When generating each distribution to be tested, DER are assigned to the feeder on a per-phase basis. For each phase, we determine the total amount of DER, $s_{Total}^{R\phi}$, as a proportion of aggregate demand:

$$s_{Total}^{R\phi} = \sum_{\forall n \in \mathcal{P}^\phi} s_n^{R\phi} = K_{DER} \sum_{\forall n \in \mathcal{P}^\phi} |p_n^{d\phi} + j q_n^{d\phi}|$$

where K_{DER}, the penetration level, is a parameter of the simulation.

The DERs on phase ϕ are then assigned to individual nodes in blocks of size s_{blk}. For each block, the node to which it is to be added is randomly selected, and blocks are sized so that it is unlikely that every node in the system has connected DER.

We create 25 DER distributions for each of three penetration levels: 135%, 120%, and 105%. For each distribution, we test both the three-phase balancing and phasor matching use cases, with objectives formulated as described in Section 3.3. In these simulations, though, the greater penetration of DER allow us to tighten the upper and lower voltage bounds of Equation (2.6), enforcing a more realistic range of 0.95 to 1.05 per unit.

Though the three-phase balancing objective remains as specified by Equation (3.7), with the shift in focus from distribution feeder to islanded system it now reflects a more critical real-world control task. We can no longer guarantee balanced operation by relying on a stiff, three-phase voltage source, and so voltage balancing must be achieved internally by controllable DERs on the feeder.

The form of the phasor matching objective is also unchanged from Equation (3.8), though for the simulations of this section we select Bus 650 to be the target node and $V_{ext} = \begin{bmatrix}1\angle 0 & 1\angle -120^\circ & 1\angle 120\end{bmatrix}^T$ as the reference to be matched. This reference value was chosen as a realistic phasor target that might be assigned in advance of the scheduled reconnection of a previously islanded microgrid to a substation.

In reporting results, the performance metric for three-phase balancing will again be the maximum voltage imbalance at a single node, as defined by the magnitude ratio of negative- and positive-sequence components: $\max_{n\in\mathcal{N}} |V_{2n}|/|V_{1n}|$. We select as our phasor matching performance metric the mismatch between the target node's phasor and the reference value, i.e. $|V_{ext} - V_{650}|$.

For a solution to be considered feasible, all nodal voltages must satisfy a convergence criterion between the LUPFM OPF and the CIM nonlinear power flow solutions. We define the convergence threshold as 10^{-5}, and it applies to both per-unit voltage magnitude and phase angle in degrees.

Extended Example 1: DER Assignment and Simulation

For illustration purposes, we describe the first of our simulation runs in detail before moving to aggregate results.

This example was carried out with a DER penetration value of 150% of the total feeder demand, i.e. $K_{DER} = 1.5$. For the IEEE13 feeder, this meant that the total DER allotted to phase A was $s^{RA}_{Total} \approx 1990$kVA, the total allotted to phase B was $s^{RB}_{Total} \approx 1850$kVA, and the total allotted to phase C was $s^{RC}_{Total} \approx 2245$kVA.

The parameter s_{blk} was set to 100kVA.

Those values resulted in the following randomized DER distribution:

Bus	632	633	634	650	651	671	675	680	684	692
s^{RA}_{Total} (kVA)	200	100	200	400	200	290.02	200	200	100	100

Table 4.1: Installed DER, phase A

Bus	632	633	634	645	646	650	651	671	675	692
s^{RB}_{Total} (kVA)	200	100	100	200	100	250.39	100	100	300	400

Table 4.2: Installed DER, phase B

Bus	632	634	645	650	651	671	680	684	692
s^{RC}_{Total} (kVA)	200	200	300	500	300	100	200	100	345.77

Table 4.3: Installed DER, phase C

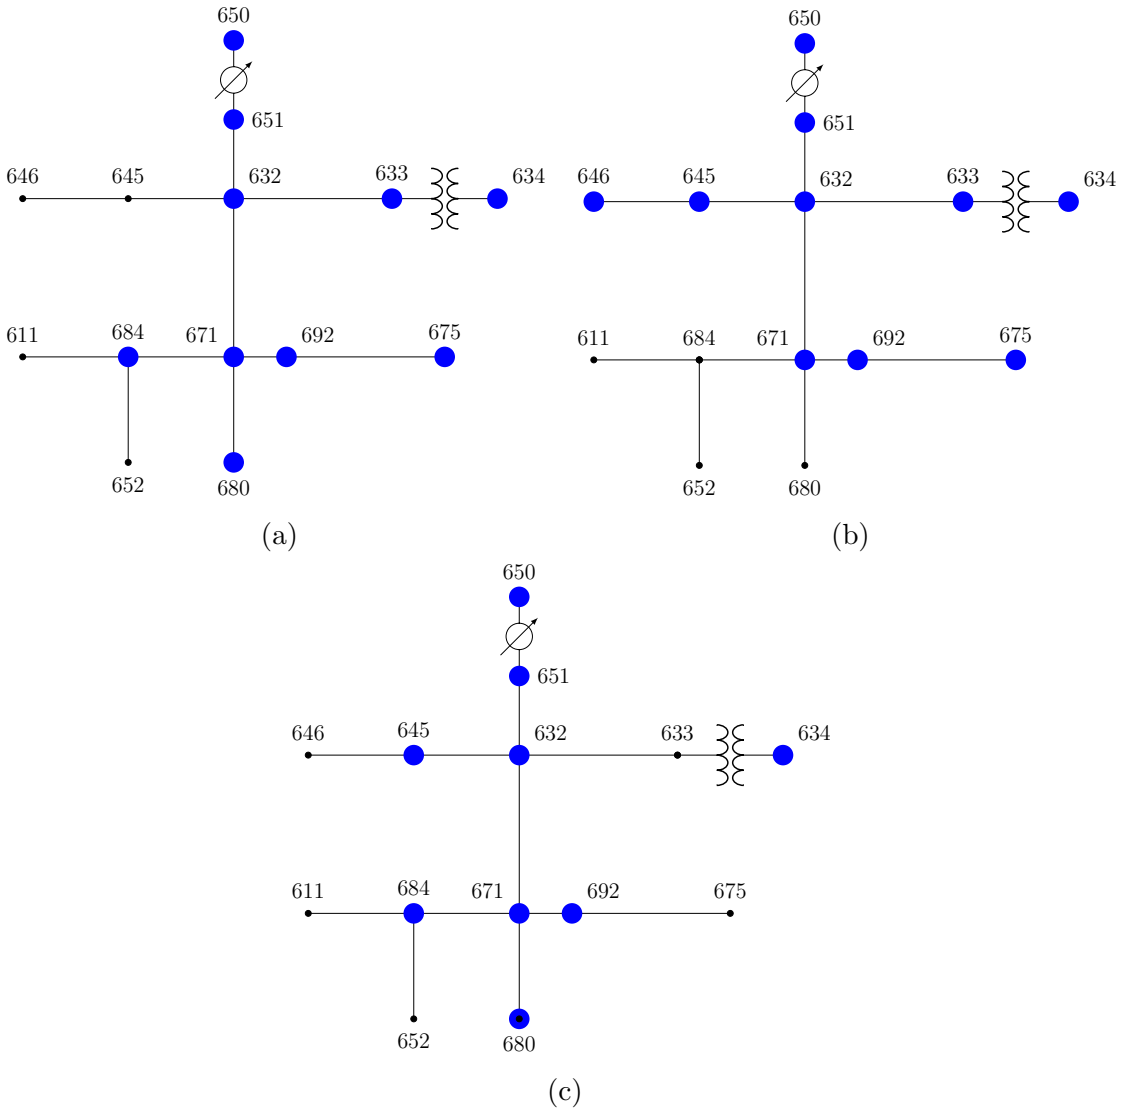

Figure 4.1: DER distribution across phases A,B, and C

During the three-phase balancing simulation, the voltage phasor profiles of the LUPFM OPF and the CIM converged to the 10^{-5} threshold after five iterations of the solution method. The most unbalanced voltage phasors on the network were those of Bus 675:

$$\mathbf{V}_{675} = \begin{bmatrix} 0.952\angle -6.84^\circ \\ 0.954\angle -126.84^\circ \\ 0.951\angle 112.84^\circ \end{bmatrix}$$

Which translates to an acceptable imbalance level of 0.26%.

During the phasor tracking simulation, the voltage phasor profiles of the LUPFM OPF and the CIM converged to the 10^{-5} threshold after three iterations of the solution method. The difference between V_{650} and V_{ext}, in terms of magnitude and angle, was

$$\epsilon_{mag} = \begin{bmatrix} 3.89 \cdot 10^{-10} \\ 1.43 \cdot 10^{-9} \\ 1.53 \cdot 10^{-10} \end{bmatrix} \qquad \epsilon_{ang} = \begin{bmatrix} 1.22 \cdot 10^{-8} \\ 6.16 \cdot 10^{-6} \\ 6.05 \cdot 10^{-6} \end{bmatrix}$$

a near-exact match.

Extended Example 2: Failure Analysis

Though the previous example highlighted a successful test case, there are combinations of DER distributions and objectives for which our solution method fails to find the required 10^{-5} agreement between LUPFM OPF and CIM. In these simulations, we have found that the method fails irrespective of the number of iterations for which the test is run. Examining one of these cases in greater detail, though, gives us a sense of the failure mechanism and the adjustments that can be made to counteract it.

The failed simulation we examine in this subsection was given a penetration level of 135%, which generated a feeder with the following DER distribution:

Bus	632	650	652	671	675	680	692
s^{RA}_{Total} (kVA)	200	400	100	291.02	400	100	300

Table 4.4: Failure case installed DER, phase A

Bus	632	633	645	646	671	675	680	692
s^{RB}_{Total} (kVA)	100	200	100	365.35	200	100	200	400

Table 4.5: Failure case installed DER, phase B

Bus	611	632	633	645	646	650	671	675	684	692
s^{RC}_{Total} (kVA)	200	200	200	300	200	100	100	100	200	421.19

Table 4.6: Failure case installed DER, phase C

We take three-phase balancing as our objective. After five iterations of the solution method, the mismatch between LUPFM OPF and CIM has settled to approximtely $8 \cdot 10^{-4}$ per unit magnitude and 0.26° in angle; too high to be useable as a PBC target. From that point, the mismatch does not change appreciably with successive iterations.

We can gain some insight into the issue by plotting the DER dispatch setpoints for phase C of Bus 646, a point where the difference between the LUPFM OPF voltage phasor and the phasor generated by the CIM was among the largest.

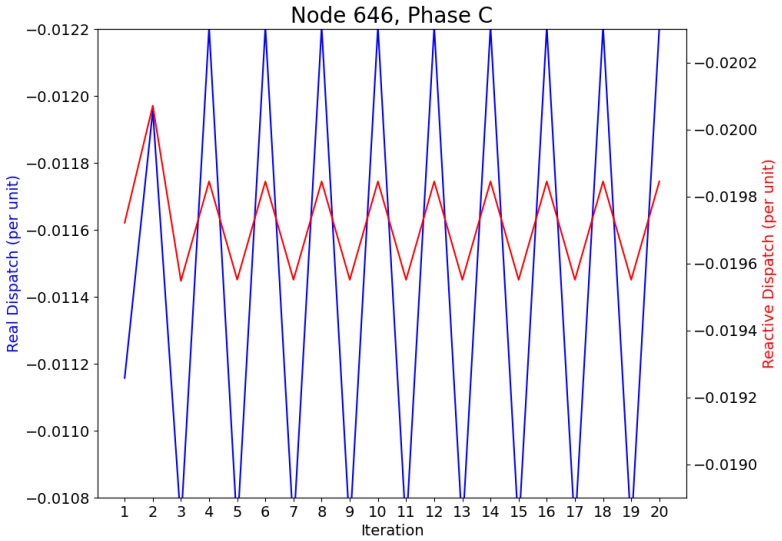

Figure 4.2: Alternation of the power dispatch setpoints of the DER at Bus 646

We can see that the power dispatch generated by the LUPFM OPF alternates between sets of values in a two-iteration cycle. This phenomenon is likely caused by circumstances similar to those of the bifurcation of a dynamical system. Here, a slight change in the estimated values of power losses and voltage phasors generated by the CIM results in an initialization of the LUPFM that returns a power dispatch qualitatively distinct from the previous iteration's. This change in dispatch then changes the CIM's nonlinear solution, which re-intializes the LUPFM in such a way that the previous iteration's solution again becomes optimal. This creates the repeating pattern of Figure 4.2 and prevents the iterative method from converging to a single, feasible solution.

We can overcome this difficulty and force the system to a single DER dispatch profile by adding a penalty factor to our objective. For three-phase balancing, this means an alteration to Equation (3.7) for all iterations after the first:
$\forall n \in \mathcal{N}; \forall nm \in \mathcal{E}; \forall \phi, \psi \in \mathcal{P}_n$:

$$\min_{\mathbf{E}_n, \boldsymbol{\theta}_n, \mathbf{p}_n^g, \mathbf{q}_n^g, \mathbf{P}_{nm}, \mathbf{Q}_{nm}} \quad \sum_{\forall n \in \mathcal{P}^3 \cup \mathcal{P}^2, \phi \neq \psi} (E_n^\phi - E_n^\psi)^2 + \sum_{\forall n \in \mathcal{P}^3 \cup \mathcal{P}^2, \phi \neq \psi} (\theta_n^\phi - \theta_n^\psi \pm \frac{2\pi}{3})^2 + \sigma_k [\sum_{\forall n \in \mathcal{N}} (\mathbf{E}_n - \mathbf{E}_{nprev})^2 + (\boldsymbol{\Theta}_n - \boldsymbol{\Theta}_{nprev})^2] \quad (4.3)$$

We add the same term to the phasor matching objective of Equation (3.8):
$\forall n \in \mathcal{N}; \forall nm \in \mathcal{E}; \forall \phi, \psi \in \mathcal{P}_n$:

$$\min_{\mathbf{E}_n, \boldsymbol{\Theta}_n, \mathbf{p}_n^g, \mathbf{q}_n^g, \mathbf{P}_{nm}, \mathbf{Q}_{nm}} \quad (\mathbf{E}_{650} - \mathbf{E}_{ext})^2 + (\boldsymbol{\Theta}_{650} - \boldsymbol{\Theta}_{ext})^2 + \sigma_k [\sum_{\forall n \in \mathcal{N}} (\mathbf{E}_n - \mathbf{E}_{nprev})^2 + (\boldsymbol{\Theta}_n - \boldsymbol{\Theta}_{nprev})^2] \quad (4.4)$$

In both Equations (4.3) and (4.4), $\mathbf{E}_{nprev}$ and $\boldsymbol{\Theta}_{nprev}$ denote the values of the voltage phasors at node n generated by the LUPFM during the iteration immediately prior. The coefficient σ_k, then, indirectly penalizes change in the LUPFM OPF's dispatch values between iterations. We define σ_k so that it increases with each iteration to ensure that the DER dispatch is forced to a fixed solution:

$$\sigma_k = 0.1 \cdot 2^{k-1} \quad (4.5)$$

where k is the index of iteration.

When the penalty term is added to our failure-case DER distribution, it results in phasor convergence between the LUPFM OPF and the CIM, though that convergence happens more slowly than was the case in previous simulations. Figure 4.3 shows the process, beginning with the third iteration so as to allow for scaling that captures the graph's features of interest.

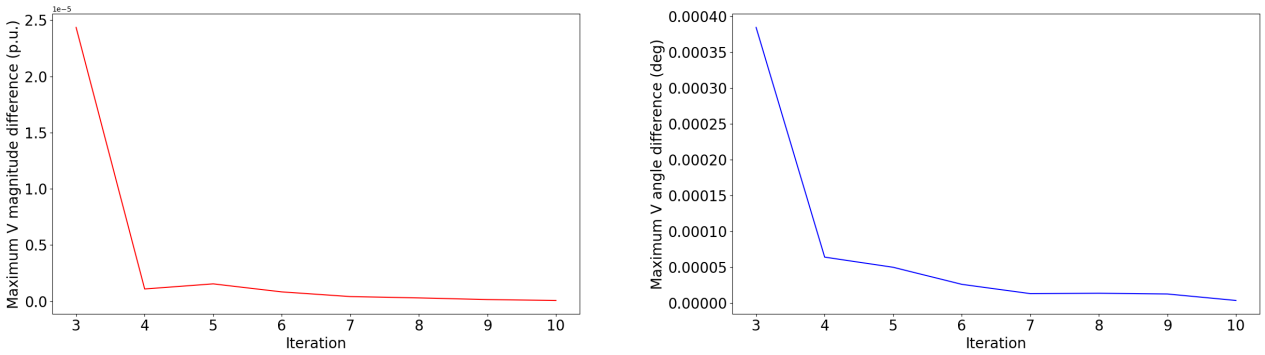

Figure 4.3: Maximum difference between LUPFM OPF and CIM phasor values with a phasor matching objective

The greatest single-node three-phase imbalance of the phasor profile generated by this simulation was 0.29%, which is comparably successful to the values of previous simulations that did not require the use of a penalty factor. That similarity in performance is not overly surprising, as the penalty factor σ_k only gains significant weight in later iterations, after the majority of the benefits of iterative optimization have already been realized.

Simulation Results in Aggregate

In this final section, we present the results of our simulations across 75 DER distributions. For each distribution, we use our iterative method to generate a phasor profile for both the three-phase balancing and phasor matching objectives specified by Equations (3.7) and (3.8). Where a solution is successfully found on the first pass through the iterative method, the value of the objective's performance metric is plotted in green.

Where the LUPFM OPF-generated solution does not converge to within 10^{-5} in agreement with the CIM within ten iterations, we re-run the simulation with the additional penalty factor of Equations (4.3) and (4.4). As earlier, we define the penalty coefficient $\sigma_k = 0.1 \cdot 2^{k-1}$, where k is the index of iteration. Where the addition of the penalty factor was able to force the system to convergence, the value of the objective's performance metric is plotted in blue.

Any simulations that did not converge even after the penalty factor was applied, or for which the LUPFM OPF failed to find a feasible dispatch, were plotted in red and the value of their performance metric marked as zero.

DER Penetration Level: 135%

At this DER penetration level, we have sufficient controllable resources to achieve both of our operational goals. In all cases across both objectives, the iterative method was able to find

a state of agreement between LUPFM OPF and CIM without the use of a penalty factor. The three-phase imbalance was always below acceptable levels, and the error in tracking the phasor target was on the order of 10^{-8} per unit magnitude and 10^{-6} degrees of angle. Without the outlier result of simulation #21, the order of the magnitude error drops to 10^{-9}.

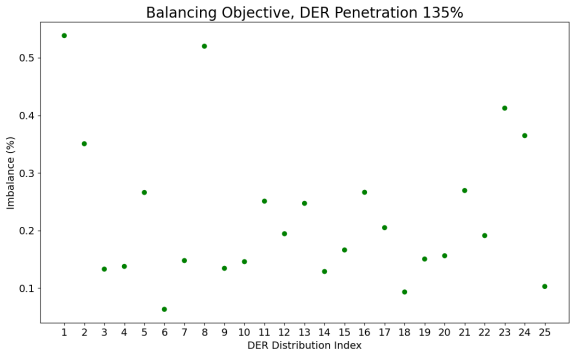

Figure 4.4: Three-phase balancing results over 25 DER distributions, 135% penetration

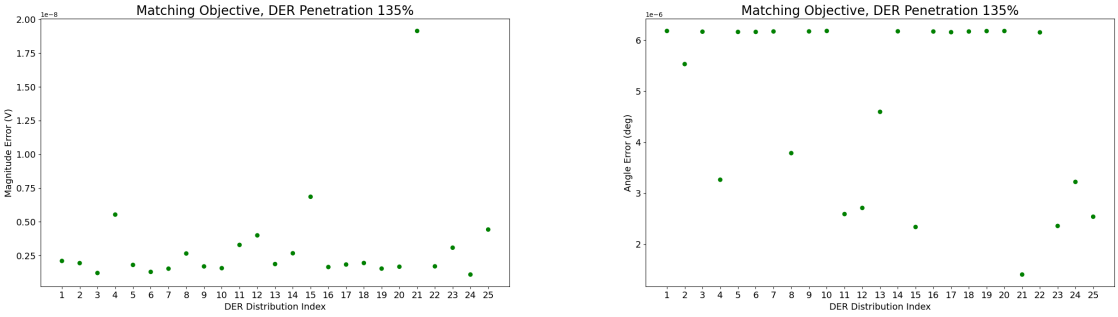

Figure 4.5: Phasor matching results over 25 DER distributions, 135% penetration

DER Penetration Level: 120%

At a penetration level of 120%, we still have sufficient controllable resources to achieve our goals without the use of a penalty factor in all test cases. The three-phase imbalance was generally higher than in the 135% case, but still acceptable. The error in tracking the phasor target was again on the order of 10^{-9} per unit magnitude and 10^{-6} degrees of angle.

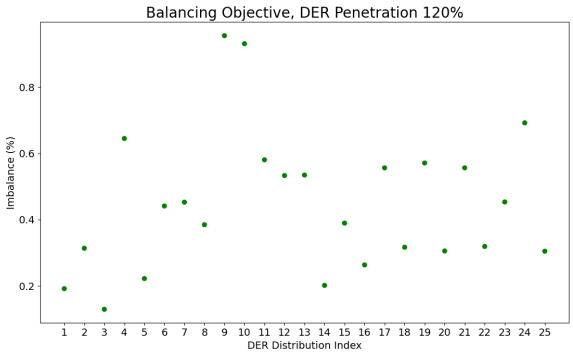

Figure 4.6: Three-phase balancing results over 25 DER distributions, 120% penetration

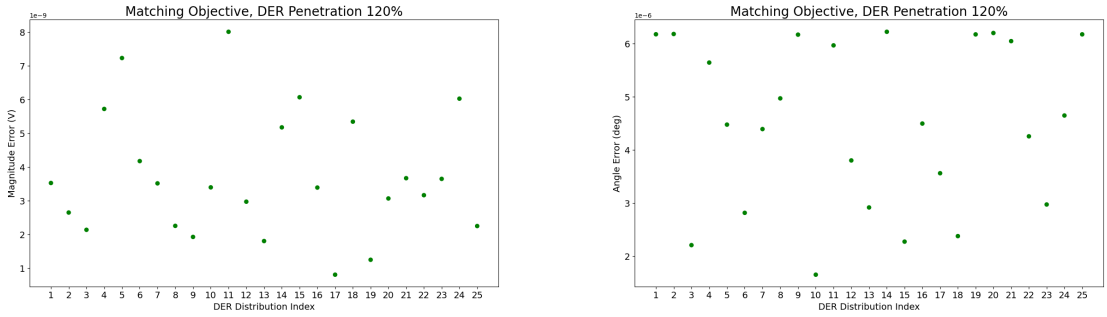

Figure 4.7: Phasor matching results over 25 DER distributions, 120% penetration

DER Penetration Level: 105%

At a penetration level of 105%, three-phase balancing was still achieved reasonably well, though in several cases the maximum imbalance at a node exceeded the 1% threshold. Phasor matching was a much more challenging prospect. Our method failed to find a solution in 14 of 25 cases, and it required a penalty factor to converge in all but one of those cases remaining.

This penetration level appears to be too low for our iterative method to be dependably used on a feeder with a random DER distribution. However, in a number of simulation runs the S-PBC did succeed in finding a feasible voltage profile that achieved its objective reasonably well. This suggests that an exploration of the relationship between the specifics of DER allocation across a feeder and the feasibility of implementing PBC could be an interesting avenue for future work, as could the effects of different QP solution techniques on the success rate of the overall method.

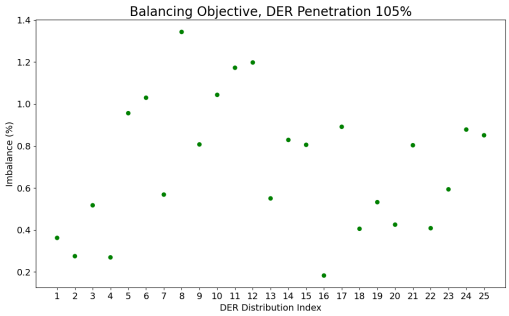

Figure 4.8: Three-phase balancing results over 25 DER distributions, 105% penetration

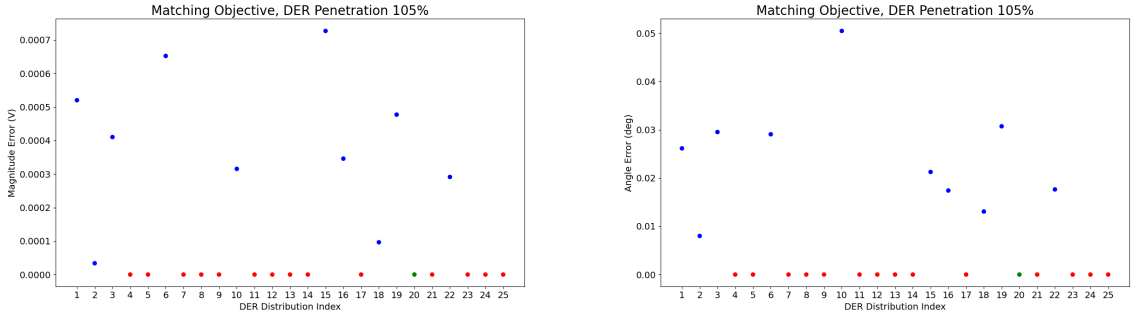

Figure 4.9: Phasor matching results over 25 DER distributions, 105% penetration

Chapter 5

Instrument Transformer Error in Distribution PMU Measurements

In this chapter, we return to a discussion of the practicalities of distribution-level PMU deployment. We focus on an important, but often overlooked, challenge that faces operational uses of phasor measurements: the error induced by instrument transformers[1]. This error has implications, not only for the PBC strategies discussed earlier in this dissertation, but for PMU applications throughout the electric grid.

We examine transformer-induced error empirically, using a nontraditional experimental test setup involving PMUs connected to an actual distribution network. This arrangement allows us a rare opportunity to track errors in situ over a time period of multiple months. After a brief introduction to instrument transformers and a motivating consideration of their calibration procedures, this chapter will end with a presentation of the data collected from our observations and a discussion of its possible effects on PBC.

While the subject matter of this chapter is very relevant to the PBC strategy, it can be considered a stand-alone analysis from the S-PBC development of Chapters 2, 3 and 4. The work presented in those earlier chapters is the primary research contribution of this dissertation.

5.1 Instrument Transformer Basics

As described in Section 2.1, modern PMUs are capable of measuring voltages and currents with extremely high accuracy. They have become an important tool for understanding the transmission grid and, as the technology has developed, distribution networks as well. Because of the high voltages involved, though, PMUs can not be connected directly to any point on either transmission or primary distribution lines. All phasor measurements taken at those levels, then, must necessarily be approximate.

[1] A portion of this work was previously submitted as a report on the ARPA-E project *Micro-Synchrophasors for Distribution Systems*, DE-AR 0000340

At the distribution level, it is possible to carry out such an approximate measurement by installing a PMU at a wall plug on the secondary side of a service transformer. However, service transformers introduce a significant amount of error, manifesting as a shift of the resulting phasor measurement by an unknown amount in both magnitude and phase angle. As a result, it is difficult to estimate the value of a voltage phasor at the primary distribution level from a secondary-side measurement without additional data and a network-level state estimator. For more accurate measurements of primary distribution voltages, operators must install specialized instrument transformers.

As a term, "instrument transformer" refers to a category that includes Potential (also called Voltage) Transformers (PTs or VTs), which are purpose-built for protecting voltage sensors while introducing a minimum of error to the measurement. Current Transformers (CTs) are another class of instrument transformer, and they introduce complications of their own, but the majority of the analysis in this chapter will focus on the PTs that are most relevant for the implementation of PBC.

PTs are generally modeled as classical non-ideal transformers, though in some cases they are coupled to capacitive voltage dividers that are included in their model as well. A diagram of a commonly accepted model of a stand-alone, non-ideal transformer [25] is reproduced in Figure 5.1.

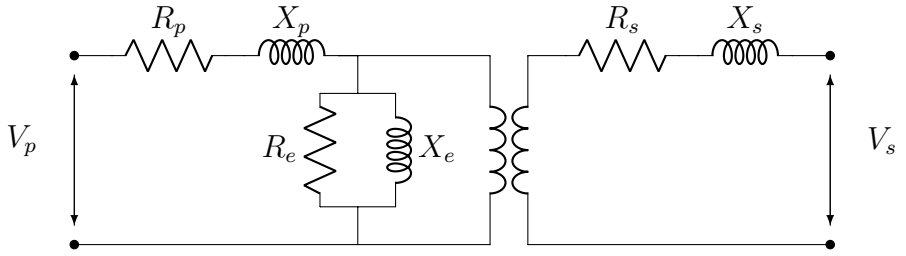

Figure 5.1: A standard model of a non-ideal transformer

Measurement error can be conceptualized as induced by the non-ideal elements of the model: the excitation impedance R_e, X_e, and the impedances of the primary and secondary windings R_p, X_p and R_s, X_s. Capacitive reactances between windings and to ground are generally large at the frequencies relevant for power systems, and they are often neglected in modeling.

Error Classes and Burdens

Instrument transformers are classified according to the error that they can be expected to induce in a measurement. The effects of that error are represented as a set of values RCF, γ, where RCF (the Ratio Correction Factor) corrects the PT's scaling of a measurement phasor's magnitude and γ corrects the PT's additive changes to its phase angle. When

describing CTs, β is generally used in place of γ to denote the phase angle error. Limits on the allowable levels of RCF and γ or β that can be induced are specified in terms of accuracy classes by IEEE Standard.

Those error limits can be graphically represented by parallelograms drawn in γ, RCF space [64]. The γ and RCF values that a transformer simultaneously presents must remain within its parallelogram in order to avoid violating the requirements of its accuracy class. An example set of PT and CT error parallelograms is reproduced in Figure 5.2.

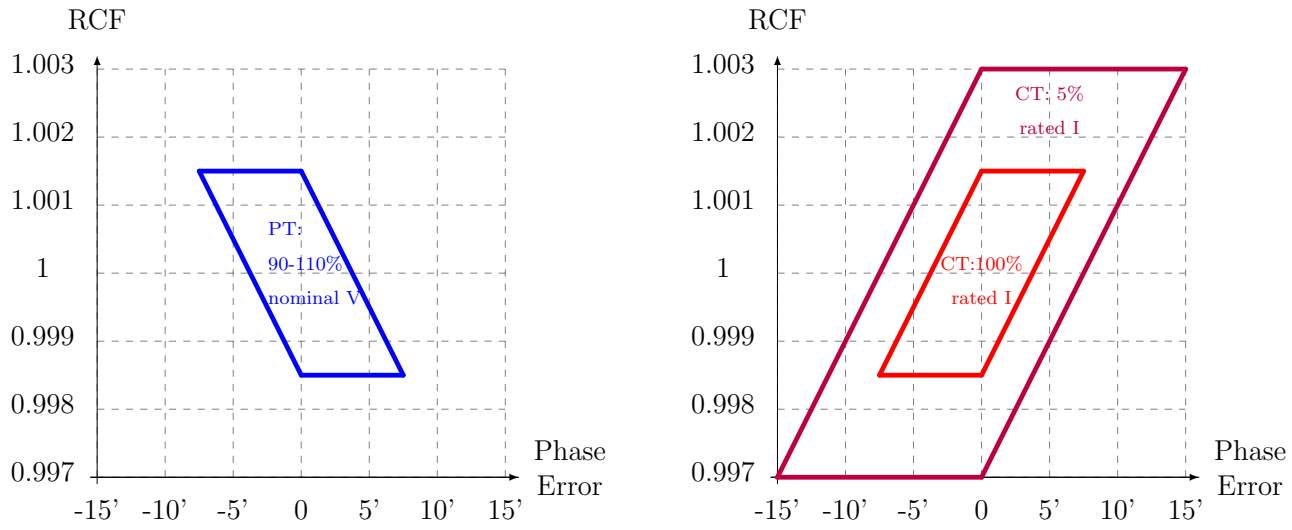

Figure 5.2: Example error parallelograms defined for a PT/CT member of the highly accurate *Class 0.15*. This figure was reproduced from [16], and based on an original from [31].

The induced error of any individual instrument transformer can be represented as a point within its error parallelogram, though that point is not known by the operator unless the transformer has been calibrated. Voltage or current data taken from the secondary side of the transformer can then be modeled as ideal phasor measurements of the primary-side phasor, multiplied by a factor of $1/RCF$ and shifted by γ or β in phase angle. This does not take into account some aspects of the PT model presented in Figure 5.1, such as the dependence of the induced error on primary-side voltage. Over a practical range of voltage values, though, that effect is relatively small [32], and the single-point representation of induced error can be considered a good approximation.

Note that the parallelograms for PTs and CTs are defined to be reflections of each other about the y-axis. This is a convention that was historically meant to ensure that PTs and CTs were built in a way which preferred that accuracy was maintained in calculations of active power. It is also worth noting that CTs experience more variation in their internal characteristics over the expected range of their primary-side measurements, and so performance requirements are relaxed at currents far below rated levels.

The definition of an error class takes into account the burden on an instrument transformer. That burden is defined as any sensors that are connected to the instrument transformer's secondary side, which draw current that affects the transformer's output in a predictable way. The burden is the key determinant of the transformer's position within its error parallelogram, and a number of standard burdens are specified by the IEEE [19].

To ensure that they stay within the IEEE's bounds over a maximum range of different burdens, PTs have slightly fewer windings on their transformer primary sides than would be expected from their nominal turns ratios. This ensures that, in an ideal case without any secondary-side burden, the PT's RCF is near its lowest allowable value. Graphically, this would be a point near the lower corner of the PT parallelogram of Figure 5.2, where the zero-burden RCF would be near 0.9985. As the burden on the transformer increases, the secondary voltage drops and the RCF increases, while γ changes according to the burden's reactance.

5.2 PT Calibration

Instrument transformer error affects measurement devices on a qualitatively small scale. For metering transformers of *Class 0.3*, the induced magnitude error is restricted to within 0.3% and the induced error in phase angle to within approximately 15 degree-minutes or 0.25°. More recently established accuracy classes, such as the *Class 0.15* whose specification is shown in Figure 5.2, have even stricter accuracy bounds. Historically, the high levels of accuracy offered by the most precise classes of instrument transformer have been sufficient for most purposes.

However, as has been noted throughout this dissertation, PBC has an exceptionally low tolerance for error in control variables. For PBC, even transformers of a highly accurate class can be problematic if they are not calibrated. As was the case in our discussion of linearization error, the impact of PT error is especially significant at the distribution level due to relatively small power flows and voltage differences between nodes. Even minor measurement errors, then, can lead to proportionately large over- or underestimates in network quantities. This is particularly true when discussing errors in the nodal voltage measurements that form a primary input to PBC.

As an aside, instrument transformer error does not affect the class of applications that rely only on measuring *variations* between successive voltage or current samples. Outside of extreme transient events, transformer-induced error is a stable characteristic that can very comfortably be considered constant over intervals on the scale of most measurement devices' sampling periods. Frequency measurement, event detection, and topology change detection are all examples of distribution PMU applications that would not be affected by instrument transformer error [77].

For an application that relies on the absolute values of voltage phasor measurements, though, PT error must be taken into account. This is done through a calibration process in which the RCF and γ values of a given PT are determined. Standard practice for

performing PT calibration involves comparing the transformer's measurement to that of a known reference, either on site or in a laboratory, and a number of variations of that technique have been suggested in the literature, e.g. [12] and [17].

There are two significant, related downsides to this calibration approach. The first is the obvious cost in time and effort. The fact that physical connections must be made to the transformer under test imposes a logistical burden on grid operators that becomes large with increasing penetration of PT-connected sensors at the distribution level.

The second concern is the length of time over which a transformer can be assumed to remain calibrated. It is known that PT parameters change as the equipment ages, as well as with variations in temperature and other local variables. For many purposes, though, PT-induced error is approximated as an invariant quantity. To the best of our knowledge, there has never been a systematic analysis of the time scale over which that assumption holds true. With the results presented in the following section, we hope to motivate that consideration.

Remote Calibration

A promising area of exploration with the potential to address the practical problems of present PT calibration methods comes in the form of remote, online techniques, many of which make use of phasor measurements. Using these methods, instrument transformer characteristics could be determined in an operational setting by analyzing voltage and current measurements throughout the network. As a strategy, this type of approach is extremely appealing.

To the extent of our knowledge, there is no field-tested method for carrying out online calibration of instrument transformers. However, we highlight one that has returned promising results in simulation. In the set of papers [52] and [53], our collaborators explore a means of using PMU measurements throughout a distribution network to simultaneously estimate line impedances and instrument transformer errors.

In both papers, the authors use the popular weighted least squares (WLS) technique on a set of PMU measurements to solve for a vector of missing network information. The first paper makes use of a linearization of Kirchhoff's current and voltage laws to carry out the WLS, and the second covers an extension of the method that allows for the treatment of networks with large deviations in line parameters. In simulation, data suggest that this method could be used to significantly reduce the measurement uncertainty associated with PT error.

In addition to [52] and [53], there have been other remote calibration methods proposed at both the transmission and distribution levels. One common thread running through many of these approaches is the use of high-quality measurements or a pre-calibrated transformer to calibrate other transformers at connected busses. Examples of this strategy can be found in [51] and [79].

Remote calibration removes a logistical burden from grid operators, and it would be an elegant means of accounting for the time variance of transformer-induced error. However,

most of the methods cited in this section require that a significant number of PMUs be installed throughout the network to be analyzed, which is presently not available at the distribution level. Though recent research efforts have provided encouraging direction for further exploration, future work will need to address this hurdle before remote calibration can become a reality.

5.3 Time Variance in PT-induced Error

In this section, we examine the degree to which it is reasonable to assume that the error induced by a PT is invariant in time.

Assuming a constant secondary-side burden, and given its dependence on the known physics of the transformer, it is certainly reasonable to assume that PT error is a slowly varying characteristic. But, as mentioned, there have been no empirical examinations of how slow that variation might be. This is, in large part, due to the inconvenience of running such a test: a redundant voltage measurement needs to be established as a truth value against which to compare measurements taken from the secondary side of a PT, and that reference needs to be maintained for the long time periods over which we might expect to see drift in that PT's characteristics.

Our team has collected the data necessary to carry out an estimation of PT error drift from a pair of μPMU-brand [45] sensors installed on a distribution feeder as part of a research deployment in Northern California. These PMUs are connected to the same single-phase bus at 7.2kV measured line-to-ground. Though the primary distribution voltage measured by both PMUs is identical, each PMU is connected through its own single-phase, *Class 0.3* PT that separately induces measurement error. This has provided us with a unique opportunity to isolate the combined effects of those PTs on the PMUs' measurements and to observe them over several years of data collection.

As described in Section 2.1, these distribution PMUs report their measurements in terms of complex phasors. The angles of those phasors are reported relative to a synchronized reference measurement taken at a nearby substation. Each PMU reported its measurements at a rate of 2 samples/second, providing us with a large dataset that can be used to establish the ratio of the errors induced by each of the two PTs.

We model the relationships of the voltage phasor measurements to the true per-unit value of the underlying primary voltage through Equations (5.1).

$$\begin{aligned} V_{1m} &= V_{true}\kappa_1 + \epsilon_1 \\ V_{2m} &= V_{true}\kappa_2 + \epsilon_2 \end{aligned} \tag{5.1}$$

where V_{true} is the true voltage value; V_{1m} and V_{2m} are the voltages measured at each of the two PMUs; κ_1 and κ_2 are the induced errors for the corresponding PTs, and ϵ_1 and ϵ_2 are the internal measurement errors of each of the two PMUs. All values are complex.

Model development for the noise introduced by the PMUs themselves was explored in detail in [9]. There, an analysis of a similarly redundant PMU deployment on the secondary

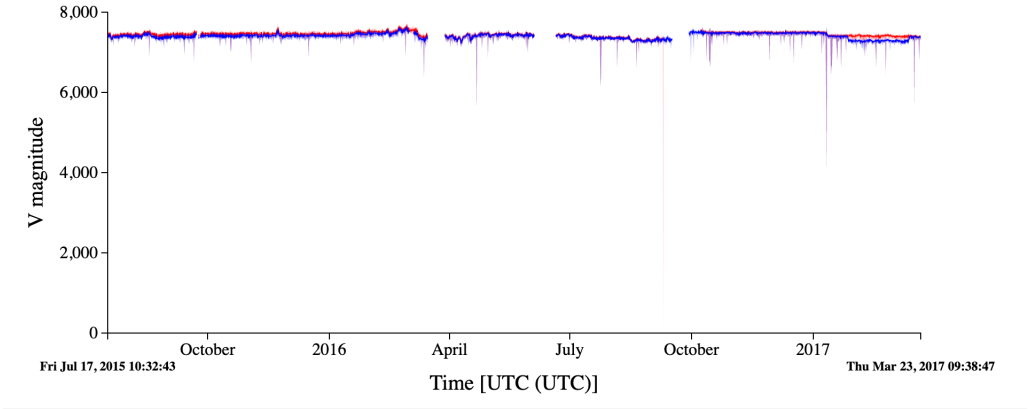

Figure 5.3: Voltage magnitude data from our paired PMUs collected over the project lifetime as a BTrDB [5] timeseries store.

distribution network found that an additive model of PMU noise with zero mean showed good correspondence to measured data. We adopt that model here in defining our ϵ_i terms, allowing us to use Ordinary Least Squares (OLS) with confidence to obtain a ratio of the two κ_i values.

We equate the V_{true} terms and rearrange Equations (5.1).

$$\begin{aligned} \frac{V_{1m} - \epsilon_1}{\kappa_1} &= \frac{V_{2m} - \epsilon_2}{\kappa_2} \\ V_{1m} &= \frac{\kappa_1}{\kappa_2} V_{2m} - \frac{\kappa_1}{\kappa_2}\epsilon_2 + \epsilon_1 \\ V_{1m} &= \kappa V_{2m} + \epsilon \end{aligned} \tag{5.2}$$

Where $\epsilon = -\frac{\kappa_1}{\kappa_2}\epsilon_2 + \epsilon_1$ remains an additive, zero-mean noise model in both real and imaginary terms, and $\kappa = \frac{\kappa_1}{\kappa_2}$. Equations (5.2) can be combined and rearranged, then solved with OLS for κ.

Results

We analyze four non-consecutive, three-day periods of data collected during 2016-2017. Each three-day dataset is separated from the next by approximately three months, ensuring that we have reported results from each of the four seasons. In the plots below, we classify them as the summer, fall, winter, and spring datasets.

For each three-day period, the data were divided into one-minute samples. The OLS performed on Equation (5.2) was used independently on each one-minute sample to fit a

value of κ. Those κ values are plotted in the figures below in terms of the magnitude of their ratio $|\kappa| = |\kappa_1|/|\kappa_2|$ and the difference between their angles $\angle\kappa = \angle\kappa_1 - \angle\kappa_2$.

Note that, on the following plots, there are fewer than the 4320 one-minute samples that would be expected for each three-day collection period. This is due to a stringent filtering requirement that was used in defining data samples: no sample was considered in the analysis if any of the 7200 data points collected over the course of the minute from either PMU was tagged with a potential error in GPS lock. This tagging is a reporting functionality provided by the μPMUs.

Summer 2016

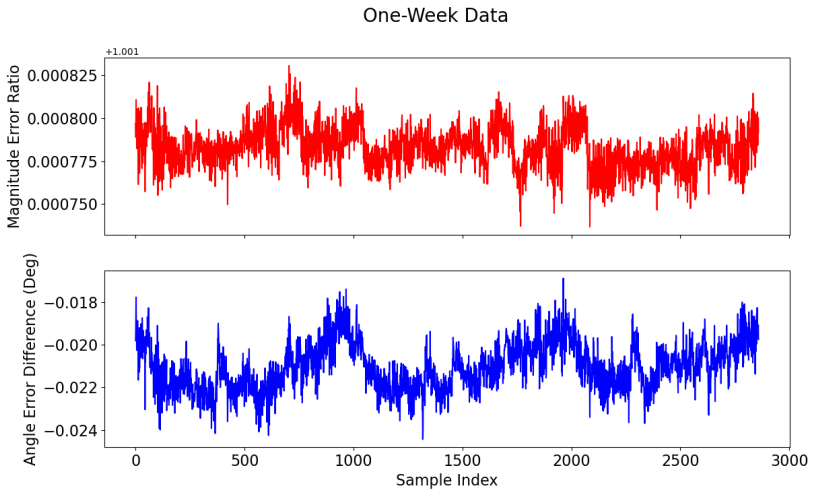

Average magnitude error ratio: 1.0018
Average angle error difference: $-0.021°$

We can see that, for these data and the three sets to follow, the change in κ within a period of several days was extremely small. Generally, the maximal shift was in the low range of the order 10^{-4} in magnitude and 10^{-2} degrees in phase angle. This agrees well with the intuition that PT error should remain approximately constant over reasonably short time periods.

As we compare the data between the nonconsecutive weeks, though, we also see a drift in the mean values of the PT error.

Fall 2016

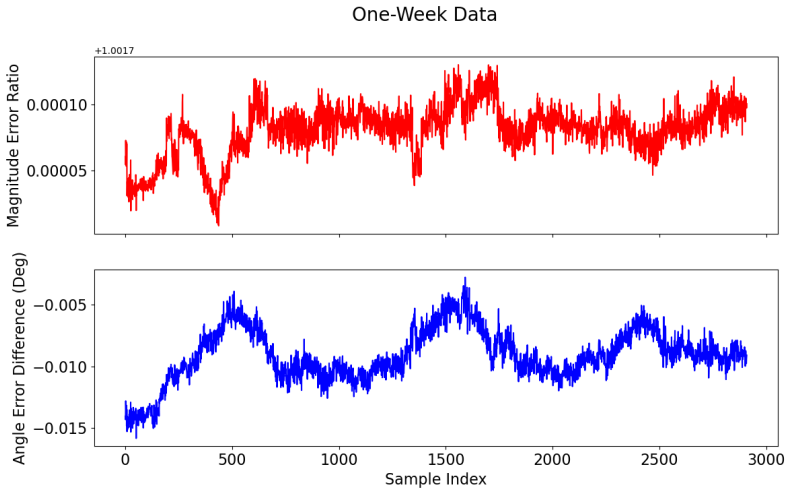

Average magnitude error ratio: 1.0018
Average angle error difference: $-0.009°$

Winter 2016/2017

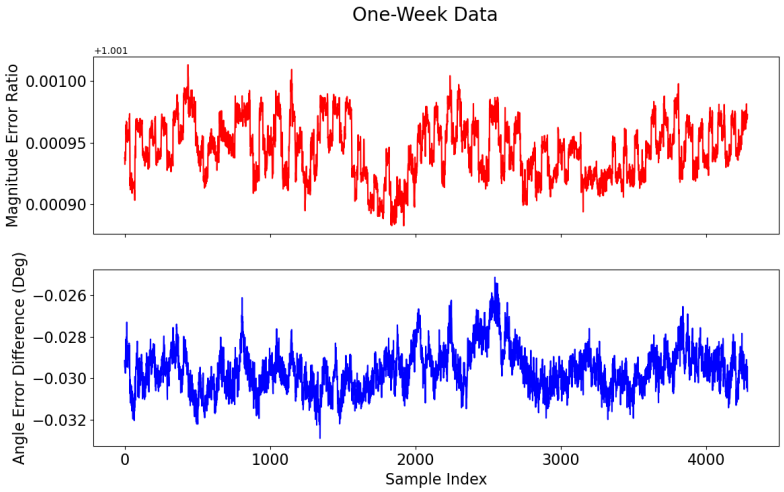

Average magnitude error ratio: 1.0019
Average angle error difference: $-0.030°$

Spring 2017

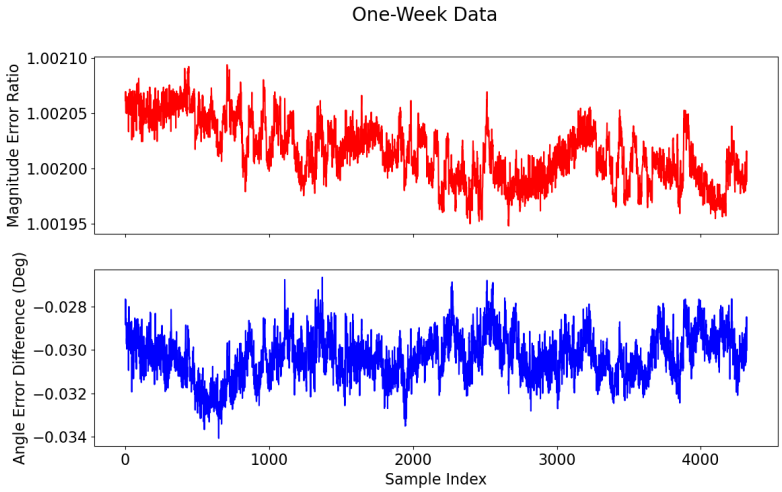

Average magnitude error ratio: 1.0020
Average angle error difference: $-0.030°$

The average of the magnitude ratio κ changes by an amount on the order of 0.01% between each of the three-month periods.

We can translate this to an approximate drift in the magnitude error induced by a single PMU by noting that, on the scale $0 \leq \Delta_{RCF} \leq 0.003$ defined by the transformer's error class, $1/(1 - \Delta RCF) \approx 1 + \Delta_{RCF}$. This means that, in a best-case scenario with minimal error induced by any one PT, each PT is responsible for half of the total change in κ.

Based on that assessment, each PT would have seen its induced error drift by approximately 0.005% over each three-month period, and 0.02% over the course of a year.

The drift in angular error showed a less steady progression, and the direction of its movement changed twice over the time period examined. The maximal drift over a three-month period was approximately $0.02°$, or 1.2 minutes of angle.

From these data, we can see that the PTs do not see a shift in their characteristics that would take them outside the bounds of their allowable error. But, projecting the observed level of parameter drift out several years, we can imagine reaching a point at which an initial calibration is no longer relevant.

This analysis is not meant to be a comprehensive treatment of transformer error. The data presented in this chapter came from a single experimental deployment and would need to be corroborated before any definitive conclusions can be drawn. However, these results do provide an important empirical touchstone, and they speak to the scale at which we might expect PT-induced measurement error to present itself.

To translate these observations into the potential effects on PBC, we return to our simple example of Chapter 3. Recall that there exists a line of impedance $|Z| \approx 0.1\Omega$ in the IEEE13 Feeder [35], and that the feeder's voltage base is 2.4kV line-to-ground.

Taking that impedance as our notional value, assuming that the RCF parameters of potential transformers on the network will drift by approximately 0.02% per year translates into a per-PT increase of roughly 10kVA erroneously dispatched by phasor-based controllers. Over time, that would become a significant source of capacity loss and general inefficiency.

Chapter 6

Conclusion

6.1 Summary of Results

This dissertation presented work that helped to define and expand PBC, a recently developed strategy for the control of DERs that harnesses emerging distribution PMU technology. Because of its unique approach, which treats the voltage phasor as the fundamental unit of decision-making and communication, PBC is exceptionally well suited to achieve a number of operational goals at the medium-voltage level.

In particular, PBC has great potential for coordinating the operation of DERs within a minigrid or a microgrid operating in island mode. In these systems, PBC is able to explicitly command power dispatch in a way that maintains balanced voltages across three phases that, without the influence of a transmission grid, are only very weakly physically coupled. In addition, PBC's phasor-target specification can facilitate microgrid reconnection, with the control of phasor angle acting to prevent an out-of-phase reclosing that could cause arcing and degrade equipment.

The PBC hierarchy defines a two-level controller, made up of the supervisory S-PBC and the local L-PBC. This work focused entirely on the S-PBC and the OPF procedures used to generate the phasor targets that it provides to lower control levels. Before attempting to extend it from standard distribution networks to islanded systems, we addressed a fundamental challenge facing our current implementation of the S-PBC's OPF component: the need to find optimal phasor profiles that are feasible solutions to the nonlinear power flow equations.

This is challenging because of our S-PBC implementation's use of the LUPFM, another recently developed control tool. The LUPFM provides a set of linearized power flow equations that treat both voltage magnitude and phase angle within its structure. It has proven to be an excellent means of generating constraints for a QP that quickly and dependably provides solutions to the OPF problem. Because the equations are linear approximations, though, there is necessarily some mismatch between their admitted solutions and sets of feasible voltage phasors.

In Chapter 3, we addressed this problem with an adapted version of a previously developed iterative solution method. The original method, introduced in [62], used successive results from a nonlinear solver to update the voltage values about which the LUPFM equations were linearized, and thus reduced the mismatch between LUPFM and nonlinear solutions at each step. We showed that, with adjustment to both the LUPFM equations and the companion nonlinear solver, the solutions generated by those two solvers could be made near-exactly equal, and therefore represent a feasible set of voltage targets. Our adjustments also enabled the generation of phasor targets on feeders that included tap-changing transformers, allowing for the application of PBC to more realistic medium-voltage networks.

We explored the extension of PBC to islanded networks in Chapter 4. We examined existing strategies for treating the slack bus, a necessary part of our chosen NR-based nonlinear power flow solver, and then designed one that proved successful in simulation. We formalized the procedure for generating S-PBC phasor targets on islanded networks, and examined one of its failure modes in detail. Based on the results of that analysis, we experimented with the addition of a penalty factor to our OPF formulation and demonstrated its effectiveness in overcoming the failure mechanism by forcing the iterative method to converge to a feasible solution with very little increase in the objective function. Finally, we defined a complete, end-to-end implementation of the S-PBC and tested it across a number of different simulated DER distributions, establishing its performance at various penetration levels.

The primary work of this dissertation was presented as a self-contained unit in Chapters 3 and 4. However, due to the importance of the issue for PBC, we briefly returned to the practical consideration of transformer-induced error in Chapter 5. This phenomenon, which is often neglected in studies of PMU applications, has the potential to affect any operations that rely on phasor measurements to be accurate in an absolute sense. In principle, any error induced by instrument transformers could be accounted for with calibration procedures. In practice, though, there will be a limited amount of time over which that calibration can be considered accurate, after which the natural change over time in those transformers' characteristics will cause the induced error to diverge from the calibrated estimation.

In Chapter 5, we took advantage of a non-traditional experimental setup to collect data from real-world PMUs on an operating distribution system. These data allowed us an empirical look at the scale of the shift that can be expected in PTs' induced errors over time. Though the data from a single deployment can not be used to draw truly rigorous conclusions, a rough calculation indicated that the drift in transformer error would translate into PBC actuation error on the order of kVA/year per PT on the network. This speaks to the need to develop a calibration strategy that could be practically carried out with regularity on systems operating under PBC.

6.2 Directions for Future Work

This work is an advancement in the development of PBC, but there are many avenues for research that must be explored before the strategy can be considered mature. A recurring

theme in PBC design has been the high level of accuracy that is necessary to provide L-PBCs with actionable phasor targets. Because the consequences of any type of error in those targets can be large when translated to the active and reactive power dispatch commanded from the network's DERs, there is much potential benefit in research efforts that focus on minimizing sources of uncertainty or calculation error in PBC implementations.

This dissertation addressed a significant source of calculation error in presenting a treatment for the mismatch introduced by the linear approximations used by our S-PBC. As mentioned, though, the development of improved calibration methods for PTs will also be necessary for the effective operation of PBC at both supervisory and local levels. Another critical line of work will be the exploration of methods for accounting for errors in feeder models, such as inaccurately recorded line impedances. Research on real-time impedance estimation methods using distribution PMUs has already shown great promise in that area, but further improvements will continue to enhance the case for PBC as a practical control strategy.

Beyond the minimization of error in phasor target generation, PBC would benefit from a systematic assessment of its performance on DER distributions with different characteristics. As we saw in Chapter 4, it was difficult for our S-PBC implementation to find feasible sets of phasor targets at low levels of DER penetration. For some of the randomized DER distributions, though, it was still possible. This suggests that the creation of tools for categorizing PBC's applicability in terms of DER siting specifics would be valuable in deployment considerations and for further development of the PBC concept.

Bibliography

[1] *A Computational Study of Preprocessed Linear Approximations to the Convex Constraints in the Iterative Linear IV ACOPF Formulation: Optimal Power Flow Paper 8*. Tech. rep. Federal Energy Regulatory Commission, June 2013.

[2] M. M. A. Abdelaziz et al. "A Novel and Generalized Three-Phase Power Flow Algorithm for Islanded Microgrids Using a Newton Trust Region Method". In: *IEEE Transactions on Power Systems* 28.1 (2013), pp. 190–201.

[3] *Accelerating the Deployment of Advanced Energy Communities: The Oakland EcoBlock*. Tech. rep. California Energy Commission, Apr. 2019.

[4] M. AlOwaifeer and A. P. S. Meliopoulos. "Centralized Microgrid Energy Management System Based on Successive Linearization". In: *2018 North American Power Symposium (NAPS)*. 2018, pp. 1–6.

[5] M. P. Andersen et al. "DISTIL: Design and implementation of a scalable synchrophasor data processing system". In: *2015 IEEE International Conference on Smart Grid Communications (SmartGridComm)*. 2015, pp. 271–277. DOI: 10.1109/SmartGridComm.2015.7436312.

[6] K. E. Antoniadou-Plytaria et al. "Distributed and Decentralized Voltage Control of Smart Distribution Networks: Models, Methods, and Future Research". In: *IEEE Transactions on Smart Grid* 8.6 (2017), pp. 2999–3008. DOI: 10.1109/TSG.2017.2679238.

[7] M. Baran and F. F. Wu. "Optimal sizing of capacitors placed on a radial distribution system". In: *IEEE Transactions on Power Delivery* 4.1 (1989), pp. 735–743.

[8] M. E. Baran and F. F. Wu. "Network reconfiguration in distribution systems for loss reduction and load balancing". In: *IEEE Transactions on Power Delivery* 4.2 (1989), pp. 1401–1407.

[9] M. Bariya, K. Moffat, and A. Von Meier. "Empirical Noise Estimation in Distribution Synchrophasor Measurements". In: *2019 International Conference on Smart Grid Synchronized Measurements and Analytics (SGSMA)*. 2019, pp. 1–7. DOI: 10.1109/SGSMA.2019.8784639.

[10] M. Baudette et al. *Hardware in the Loop Benchmarking for Phasor-Based Control Validation (in preparation)*.

[11] A. Bonfiglio et al. "Modeling and Experimental Validation of an Islanded No-Inertia Microgrid Site". In: *IEEE Transactions on Sustainable Energy* 9.4 (2018), pp. 1812–1821.

[12] A. Brandolini, M. Faifer, and R. Ottoboni. "A Simple Method for the Calibration of Traditional and Electronic Measurement Current and Voltage Transformers". In: *IEEE Transactions on Instrumentation and Measurement* 58.5 (2009), pp. 1345–1353. DOI: 10.1109/TIM.2008.2009184.

[13] K. Bunker, K. Hawley, and J. Morris. *Renewable Microgrids: Profiles from Islands and Remote Communities Across the Globe*. Rocky Mountain Institute, Nov. 2015.

[14] E. J. Coster et al. "Integration Issues of Distributed Generation in Distribution Grids". In: *Proceedings of the IEEE* 99.1 (2011), pp. 28–39. DOI: 10.1109/JPROC.2010.2052776.

[15] S. Dhople et al. "Reexamining the Distributed Slack Bus". In: *IEEE Transactions on Power Systems* (2020), pp. 1–1. DOI: 10.1109/TPWRS.2020.2987325.

[16] *DisTT: Synchrophasor Monitoring for Distribution Systems-Technical Foundations and Applications*. Tech. rep. NASPI-2018-TR-001. North American Synchrophasor Initiative, 2018. DOI: 10.1109/IEEESTD.2005.97975.

[17] K. Draxler and R. Styblikova. "Calibrating an Instrument Voltage Transformer to achieve reduced uncertainty". In: *2009 IEEE Instrumentation and Measurement Technology Conference*. 2009, pp. 1556–1561. DOI: 10.1109/IMTC.2009.5168703.

[18] G. Díaz, J. Gómez-Aleixandre, and J. Coto. "Direct Backward/Forward Sweep Algorithm for Solving Load Power Flows in AC Droop-Regulated Microgrids". In: *IEEE Transactions on Smart Grid* 7.5 (2016), pp. 2208–2217. DOI: 10.1109/TSG.2015.2478278.

[19] G.E. Digital Energy. *Instrument Transformer Technical Information and Application*. Dec. 2009. URL: https://www.gegridsolutions.com/iti/specs/resources.htm.

[20] M.N. Faqiry et al. "Transactive-Market-Based Operation of Distributed Electrical Energy Storage with Grid Constraints". In: *Energies* 10 (2017), p. 1891.

[21] *Final Report on the August 14, 2003 Blackout in the United States and Canada: Causes and Recommendations*. Tech. rep. U.S.-Canada Power System Outage Task Force, Apr. 2004.

[22] C. L. Fortescue. "Method of Symmetrical Co-Ordinates Applied to the Solution of Polyphase Networks". In: *Transactions of the American Institute of Electrical Engineers* XXXVII.2 (1918), pp. 1027–1140.

[23] L. Gan and S. H. Low. "Convex relaxations and linear approximation for optimal power flow in multiphase radial networks". In: *2014 Power Systems Computation Conference*. 2014, pp. 1–9.

[24] P. A. N. Garcia et al. "Three-phase power flow calculations using the current injection method". In: *IEEE Transactions on Power Systems* 15.2 (2000), pp. 508–514.

[25] J.J. Grainger and W.D. Stevenson. *Power system analysis*. McGraw-Hill series in electrical and computer engineering: Power and energy. McGraw-Hill, 1994. ISBN: 9780070612938. URL: https://books.google.com/books?id=NBIoAQAAMAAJ.

[26] A. Hampson et al. *Combined Heat and Power (CHP) Technical Potential in the United States*. Tech. rep. DOE/EE-1328. U.S. Department of Energy Office of Energy Efficiency and Renewable Energy, Mar. 2016.

[27] Adam Hirsch, Yael Parag, and Josep Guerrero. "Microgrids: A review of technologies, key drivers, and outstanding issues". English. In: *Renewable and Sustainable Energy Reviews* 90 (July 2018), pp. 402–411. ISSN: 1364-0321. DOI: 10.1016/j.rser.2018.03.040.

[28] *History of Optimal Power Flow and Formulations: Optimal Power Flow Paper 1*. Tech. rep. Federal Energy Regulatory Commission, Dec. 2012.

[29] G. P. Holdmann, R. W. Wies, and J. B. Vandermeer. "Renewable Energy Integration in Alaska's Remote Islanded Microgrids: Economic Drivers, Technical Strategies, Technological Niche Development, and Policy Implications". In: *Proceedings of the IEEE* 107.9 (2019), pp. 1820–1837.

[30] K. Horowitz et al. *An Overview of Distributed Energy Resource (DER) Interconnection: Current Practices and Emerging Solutions*. Tech. rep. NREL/TP-6A20-72102. National Renewable Energy Laboratory, Apr. 2019.

[31] "IEEE Standard for High-Accuracy Instrument Transformers". In: *IEEE Std C57.13.6-2005* (2005). DOI: 10.1109/IEEESTD.2005.97975.

[32] *Instrument Transformers Application Guide*. Tech. rep. 1HSM 9543 40-00en, Edition 4. ABB, 2015.

[33] S. A. U. Islam et al. "Phasor-Based Adaptive Control of a Test-Feeder Distribution Network: Application of Retrospective Cost Adaptive Control to the IEEE 13-Node Test Feeder". In: *IEEE Control Systems Magazine* 39.4 (2019), pp. 56–74.

[34] M. Z. Kamh and R. Iravani. "Unbalanced Model and Power-Flow Analysis of Microgrids and Active Distribution Systems". In: *IEEE Transactions on Power Delivery* 25.4 (2010), pp. 2851–2858.

[35] W. H. Kersting. "Radial distribution test feeders". In: *2001 IEEE Power Engineering Society Winter Meeting. Conference Proceedings (Cat. No.01CH37194)*. Vol. 2. 2001, 908–912 vol.2.

[36] A. Kulmala, S. Repo, and P. Järventausta. "Coordinated Voltage Control in Distribution Networks Including Several Distributed Energy Resources". In: *IEEE Transactions on Smart Grid* 5.4 (2014), pp. 2010–2020. DOI: 10.1109/TSG.2014.2297971.

[37] P. Kundur, N.J. Balu, and M.G. Lauby. *Power System Stability and Control*. EPRI power system engineering series. McGraw-Hill Education, 1994. ISBN: 9780070359581. URL: https://books.google.com/books?id=2cbvyf8Ly4AC.

[38] R. H. Lasseter. "Smart Distribution: Coupled Microgrids". In: *Proceedings of the IEEE* 99.6 (2011), pp. 1074–1082.

[39] R.H. Lasseter et al. *Integration of Distributed Energy Resources: The CERTS MicroGrid Concept*. Tech. rep. LBNL-50829. Consortium for Electric Reliability Technology Solutions, 2003.

[40] R. Li, Q. Wu, and S. S. Oren. "Distribution Locational Marginal Pricing for Optimal Electric Vehicle Charging Management". In: *IEEE Transactions on Power Systems* 29.1 (2014), pp. 203–211.

[41] S. H. Low. "Convex Relaxation of Optimal Power Flow—Part II: Exactness". In: *IEEE Transactions on Control of Network Systems* 1.2 (2014), pp. 177–189.

[42] R. McAllister et al. *New Approaches to Distributed PV Interconnection: Implementation Considerations for Addressing Emerging Issues*. Tech. rep. NREL/TP-6A20-72038. National Renewable Energy Laboratory, 2019.

[43] A. von Meier. *Electric Power Systems: A Conceptual Introduction*. Wiley Survival Guides in Engineering and Science. Wiley, 2006. ISBN: 9780470036402. URL: https://books.google.com/books?id=bWAi22IB3lkC.

[44] M. Michel. "Microgrids for Mines: Global Perspectives". In: *Energy and Mines World Congress*. Nov. 2017. URL: https://essinc.com/2017/11/16/microgrids-mines-global-perspectives/.

[45] *microPMU datasheet*. Power Standards Lab, 2019. URL: https://www.powerstandards.com/product/micropmu/highlights/.

[46] *Mini Grids for Half a Billion People : Market Outlook and Handbook for Decision Makers*. Tech. rep. 014/19. World Bank, Energy Sector Management Assistance Program, June 2019.

[47] F. Mumtaz et al. "A Novel Approach to Solve Power Flow for Islanded Microgrids Using Modified Newton Raphson With Droop Control of DG". In: *IEEE Transactions on Sustainable Energy* 7.2 (2016), pp. 493–503.

[48] Mohammed E. Nassar and M.M.A. Salama. "A novel branch-based power flow algorithm for islanded AC microgrids". In: *Electric Power Systems Research* 146 (2017), pp. 51 –62. ISSN: 0378-7796. DOI: https://doi.org/10.1016/j.epsr.2017.01.019. URL: http://www.sciencedirect.com/science/article/pii/S0378779617300287.

[49] H. T. Nguyen and L. B. Le. "Optimal energy management for building microgrid with constrained renewable energy utilization". In: *2014 IEEE International Conference on Smart Grid Communications (SmartGridComm)*. 2014, pp. 133–138.

[50] H. Nikkhajoei and R. Iravani. "Steady-State Model and Power Flow Analysis of Electronically-Coupled Distributed Resource Units". In: *IEEE Transactions on Power Delivery* 22.1 (2007), pp. 721–728.

[51] A. Pal et al. "Online Calibration of Voltage Transformers Using Synchrophasor Measurements". In: *IEEE Transactions on Power Delivery* 31.1 (2016), pp. 370–380. DOI: 10.1109/TPWRD.2015.2494058.

[52] P.A. Pegoraro et al. "Handling Instrument Transformers and PMU Errors for the Estimation of Line Parameters in Distribution Grids". In: *IEEE International Workshop on Applied Measurmeents for Power Systems (AMPS)* (2017).

[53] P.A. Pegoraro et al. "Line Impedance Estimation based on Synchrophasor Measurements for Power Distribution Systems". In: *IEEE Transactions on Instrumentation & Measurement [In preparation]* (2018).

[54] I. Pérez-Arriaga, C. Knittle, et al. *Utility of the Future: An MIT Energy Initiative Response to an Industry in Transition*. Tech. rep. MIT Energy Initiative, 2016.

[55] A. G. Phadke. "Synchronized phasor measurements-a historical overview". In: *IEEE/PES Transmission and Distribution Conference and Exhibition*. Vol. 1. 2002, 476–479 vol.1.

[56] A. G. Phadke. "Synchronized phasor measurements in power systems". In: *IEEE Computer Applications in Power* 6.2 (1993), pp. 10–15.

[57] A. G. Phadke and J. S. Thorp. "HISTORY AND APPLICATIONS OF PHASOR MEASUREMENTS". In: *2006 IEEE PES Power Systems Conference and Exposition*. 2006, pp. 331–335.

[58] P. Piagi and R. H. Lasseter. "Autonomous control of microgrids". In: *2006 IEEE Power Engineering Society General Meeting*. 2006.

[59] *Renewables 2020 Executive Summary*. Tech. rep. International Energy Agency, 2020.

[60] A. D. Rodriguez, F. M. Fuentes, and A. J. Matta. "Comparative analysis between voltage unbalance definitions". In: *2015 Workshop on Engineering Applications - International Congress on Engineering (WEA)*. 2015, pp. 1–7.

[61] S. R. Sanders et al. "Generalized averaging method for power conversion circuits". In: *IEEE Transactions on Power Electronics* 6.2 (1991), pp. 251–259.

[62] M. Sankur. "Optimal Control of Commercial Office Battery Systems, and Grid Integrated Energy Resources on Distribution Networks". PhD thesis. UC Berkeley, 2017.

[63] M. D. Sankur et al. "Model-Free Optimal Voltage Phasor Regulation in Unbalanced Distribution Systems". In: *IEEE Transactions on Smart Grid* 11.1 (2020), pp. 884–894.

[64] J. L. Settles, W. R. Farber, and E. E. Conner. "The Analytical and Graphical Determination of Complete Potential Transformer Characteristics". In: *Transactions of the American Institute of Electrical Engineers. Part III: Power Apparatus and Systems* 79.3 (1960), pp. 1213–1218. DOI: 10.1109/AIEEPAS.1960.4500942.

[65] Shiqiong Tong, M. Kleinberg, and K. Miu. "A distributed slack bus model and its impact on distribution system application techniques". In: *2005 IEEE International Symposium on Circuits and Systems*. 2005, 4743–4746 Vol. 5.

[66] J. W. Smith et al. "Smart inverter volt/var control functions for high penetration of PV on distribution systems". In: *2011 IEEE/PES Power Systems Conference and Exposition*. 2011, pp. 1–6.

[67] A. M. Stankovic, S. R. Sanders, and T. Aydin. "Dynamic phasors in modeling and analysis of unbalanced polyphase AC machines". In: *IEEE Transactions on Energy Conversion* 17.1 (2002), pp. 107–113.

[68] E. M. Stewart et al. "Accuracy and validation of measured and modeled data for distributed PV interconnection and control". In: *2015 IEEE Power & Energy Society General Meeting*. 2015, pp. 1–5.

[69] Emma Stewart et al. *Analysis of High-Penetration Levels of Photovoltaics into the Distribution Grid on Oahu, Hawaii: Detailed Analysis of HECO Feeder WF1*. Tech. rep. NREL/SR-5500-54494. National Renewable Energy Lab. (NREL), May 2013.

[70] *Technology transition final public report: smart power infrastructure demonstration for energy reliability and security (SPIDERS)*. Tech. rep. Naval Facilities Engineering Command; U.S. Department of Defense, Dec. 2015.

[71] W. F. Tinney and C. E. Hart. "Power Flow Solution by Newton's Method". In: *IEEE Transactions on Power Apparatus and Systems* PAS-86.11 (1967), pp. 1449–1460.

[72] Dan T. Ton and Merrill A. Smith. "The U.S. Department of Energy's Microgrid Initiative". In: *The Electricity Journal* 25.8 (2012), pp. 84 –94. ISSN: 1040-6190. DOI: https://doi.org/10.1016/j.tej.2012.09.013. URL: http://www.sciencedirect.com/science/article/pii/S1040619012002254.

[73] *Tracking SDG7: The Energy Progress Report*. Tech. rep. International Energy Agency, 2020.

[74] *United States Electricity Industry Primer*. Tech. rep. DOE/OE-0017. U.S. Department of Energy Office of Electricity Delivery and Energy Reliability, July 2015.

[75] *U.S. Energy Storage Monitor Executive Summary*. Tech. rep. Wood Mackenzie Power and Renewables, June 2020.

[76] A. von Meier et al. "Phasor-Based Control for Scalable Integration of Variable Energy Resources". In: *Energies* 13.1 (2020), p. 190.

[77] A. von Meier et al. "Precision Micro-Synchrophasors for Distribution Systems: A Summary of Applications". In: *IEEE Transactions on Smart Grid* 8.6 (2017), pp. 2926–2936.

[78] B. Washom et al. "Ivory Tower of Power: Microgrid Implementation at the University of California, San Diego". In: *IEEE Power and Energy Magazine* 11.4 (2013), pp. 28–32.

[79] Z. Wu et al. "Three-phase instrument transformer calibration with synchronized phasor measurements". In: *2012 IEEE PES Innovative Smart Grid Technologies (ISGT)*. 2012, pp. 1–6. DOI: 10.1109/ISGT.2012.6175540.