

Article

Phasor-Based Control for Scalable Integration of Variable Energy Resources

Alexandra von Meier ^{1,*}, Elizabeth L. Ratnam ², Kyle Brady ¹, Keith Moffat ¹ and Jaimie Swartz ¹

¹ Department of Electrical Engineering, University of California, Berkeley, CA 94720, USA; kwbrady@berkeley.edu (K.B.); keithm@berkeley.edu (K.M.); jaimie.swartz@berkeley.edu (J.S.)

² Research School of Electrical, Energy and Materials Engineering, Australian National University, Canberra, ACT 2601, Australia; elizabeth.ratnam@anu.edu.au

* Correspondence: vonmeier@berkeley.edu

† Current address: 253 Cory Hall, Berkeley, CA 94720, USA.

Received: 28 November 2019; Accepted: 26 December 2019; Published: 1 January 2020



Abstract: We propose an innovative framework termed phasor-based control (PBC) to facilitate the integration of heterogeneous and intermittent distributed energy resources (DER) on the electric grid. PBC presents a unified approach that is agnostic to optimization criteria and to the particular characteristics of participating resources. It is enabled by synchronized, high-precision voltage phasor measurements that allow stating control objectives in grid-specific, rather than resource-specific, terms. We present qualitative justification and examine the general feasibility of this control approach, including the behavior of candidate control algorithms in simulation. Initial results suggest that PBC has significant potential to support stable and resilient grid operations in the presence of arbitrarily high penetrations of DER.

Keywords: synchrophasors; distributed energy resources; phasor-based control; electric power grids; renewable power integration

1. Introduction

Electric power generation from renewable, variable, and distributed resources is increasing dramatically in many countries [1]. Other cost-effective technologies, such as all-electric vehicles, automated demand response, and customer- or community-scale energy storage, are also poised for significant growth. These distributed energy resources (DERs) introduce both challenges and opportunities at the periphery of the grid. Specific issues of concern with residential-scale solar photovoltaics (PV) include voltage rise or protection coordination problems when PV generation exceeds co-located load [2,3], and network congestion due to solar variability at times of peak demand [4]. For these reasons, distribution utilities may limit the connection (or permissible levels) of PV generation, e.g., by defining feeder hosting capacities. Electric vehicle charging, if not carefully controlled, may similarly strain the limits of legacy distribution systems [5]. Local energy storage, such as customer- or utility-owned batteries, potentially mitigate variability, but control objectives can be conflicting [4]. Notably, resources may create local problems at the distribution level, while acting to solve system-wide problems at the transmission level: for example, causing voltage volatility, while providing frequency regulation services to the bulk grid. To support the cost-effective integration of environmentally desirable technologies and to assist grid operators in delivering reliable, high-quality power, more sophisticated approaches for controlling heterogeneous DER are needed [6–11].

In this paper, we propose a unified framework termed phasor-based control (PBC) to facilitate dramatic increases in renewable generation levels throughout the electric grid, including the

distribution network. PBC frames the contributions of generators and loads in terms of their physical effect on the network, by explicitly referring to the electrical state variables, voltage magnitude V and voltage phase angle δ , that constitute the voltage *phasor* at any given network node. (The phasor representation assumes the voltage is a sinusoid of the form $v(t) = V_{max} \cos(\omega t + \delta) = \text{Re}\{V_{max} e^{j\omega t} e^{j\delta}\}$, where ω is the (constant) angular frequency $2\pi 50$ or $2\pi 60$, and $j = \sqrt{-1}$. In power engineering convention, we ignore the $e^{j\omega t}$ term and write $V_{max} e^{j\delta}$ as $V \angle \delta$, where $V = \frac{1}{\sqrt{2}} V_{max}$ is the root-mean-square magnitude.)

Our proposed PBC approach only recently became feasible with the development of ultra-precise synchronized phasor measurement units (μ PMUs) that provide synchrophasor data with sufficient resolution for the power distribution context: instruments can now reliably discern angle differences as small as ten millidegrees, or about half a microsecond [12,13]. In controlling resources directly based on physical measurements from the grid, PBC fundamentally differs from approaches that aim to manage resources through refined economic methods, such as time-varying locational prices. To our knowledge, nowhere has a comprehensive, integrated PBC scheme been proposed to encompass real and reactive power flow control across transmission and distribution systems.

The problem of controlling large numbers of distributed resources to help rather than hinder electric grid operation is well recognized in both academic literature and industry. To date, most emphasis has been on the temporal coordination of diverse resources to satisfy overall power balance in the grid. Key innovations in this area include the refinement of time-varying price signals toward faster time steps [14], automated demand response [15], and eligibility of small-scale resources to participate in ancillary service markets, such as frequency regulation [16]. The increasing participation of distributed resources in grid management, however, comes with a need to account for the spatial dimension, given the highly location-specific constraints of distribution networks that must still be satisfied while optimizing on a larger geographic scale.

The optimal recruitment of distributed energy storage resources is a case in point. Here, a single resource can allocate portions of its power and energy capacity to simultaneously provide multiple services on different time scales; for example, energy arbitrage and frequency regulation [17]. The resource may also serve competing purposes, e.g., peak shaving for the utility versus backup power reserve for the customer in case of an outage [18]. However, to maximize the overall benefits of the resource, it is necessary to account for its particular local impact on the network infrastructure, which can be negative or positive: namely, does the resource create or does it mitigate excessive voltage drops, thermal overload risks, or reversal of power flow direction on the distribution feeder?

One general solution strategy is to extend well-understood optimization tools from the transmission into the distribution realm. For example, Locational Marginal Pricing (LMP) could be applied to capture distribution congestion and discriminate the value of power injections with higher granularity in both time and space [19]. Note that to properly account for differential physical impacts on distribution circuits, LMPs would have to be assigned on a small scale, potentially varying from one customer to the next (e.g., depending on distance from the substation) in a way that is not transparent to the ratepayer.

Any pricing approach will rely on the use of system models and some form of centralized optimization. Specifically, knowledge of the grid and its operating state may be converted into a suitable power or energy price at a given location and time. If that price is to reflect local as well as global conditions, it must be computed with information from both levels. The more centralized a coordination strategy, the greater a data acquisition and communications burden it imposes on grid operations, along with the need for better network model information (i.e., connectivity and line impedances). However, model data at the distribution level is often incomplete or less precise than transmission level data, and a dependence on too much information introduces an inherent vulnerability into any control scheme.

To reduce dependence on network model information, autonomous or decentralized and model-free control strategies are an attractive alternative or complement to centralized optimization

where possible. For example, the highly location-sensitive voltage regulation function of distributed generators can be decoupled from power and energy provision by controlling reactive power output on a volt-VAR droop curve. Such a curve assigns a specific amount of reactive power injection as a function of, locally measured voltage magnitude [20,21]. In this way, while real power injection may follow some other criteria such as price-based optimization, positive or negative reactive power injection can serve to mitigate local network constraints without the need for centralized control. It has also been demonstrated that active control of generation resources based on a voltage phase angle separation constraint on a distribution line is effective at limiting real power flow [22].

Although controlling real and reactive power separately is a practical and sensible near-term solution to feeder voltage problems [23], this decoupled approach has fundamental limitations in the long term. Modulating real power is necessary to navigate distribution system constraints, such as capacity limits or protection coordination under reverse (real) power flow. Furthermore, at high DER penetration levels, managing voltage magnitude using only reactive power can yield unsatisfactory results with excessive losses if real and reactive power flow in opposite directions [24]. An underlying difficulty is that distribution circuits conform poorly to the standard assumption that reactive power Q varies mainly with the voltage magnitude V , while real or active power P varies mainly with voltage phase angle δ . Decoupling applies when inductive reactance X is much greater than resistance R , as is generally true in transmission systems. Since distribution networks have smaller physical dimensions, it follows that distribution feeders tend to have significant resistive components, and P and Q are consequently intertwined with both V and δ . This inter-relationship is most clearly seen in a pair of equations based on the DistFlow equations [25] adapted by [24,26] for the voltage magnitudes and angles at adjacent nodes 1 and 2:

$$|V_1|^2 - |V_2|^2 = 2(RP + XQ) + (R^2 + X^2) \frac{(P^2 + Q^2)}{|V_1|^2}, \quad (1)$$

$$\sin(\delta_1 - \delta_2) = \frac{XP - RQ}{|V_1||V_2|}, \quad (2)$$

where P and Q flow from node 1 to 2, and $Z = R + jX$ is the impedance of the connecting circuit branch. The unbalanced three-phase formulation derived in [24] expands into vector and matrix quantities but is otherwise similar. These equations can be approximated by eliminating the loss term in Equation (1) and applying the small angle approximation to Equation (2):

$$|V_1|^2 - |V_2|^2 \approx 2(RP + XQ), \quad (3)$$

$$\delta_1 - \delta_2 \approx \frac{XP - RQ}{|V_1||V_2|}. \quad (4)$$

These equations can be further approximated by the linear equations:

$$|V_1| - |V_2| \approx RP + XQ, \quad (5)$$

$$\delta_1 - \delta_2 \approx XP - RQ. \quad (6)$$

The interdependency of P , Q and V , δ in distribution circuits motivates our integrated PBC framework, where real and reactive power are modulated based on the entire measured voltage phasor. Beyond yielding intuitive insight, the above linearization also yields a convex power flow equality, which is leveraged for the supervisory PBC optimization problem [27].

2. The PBC Paradigm: Motivating Examples

In this section, we present toy examples from the distribution and transmission context that motivate and validate the PBC paradigm, with basic simulation results in MATLAB and ePHASORSim.

The ePHASORSim simulation package from OPAL-RT integrates quasi steady-state power analysis with control logic developed in Simulink, a MATLAB-based graphical programming environment. We will refer readers to other papers from our extended project team for more detailed control algorithms and simulation results.

PBC computes power injection set-points throughout the a.c. network by tracking the state variables V and δ that uniquely determine real and reactive power flows. In our hierarchical framework, one or more supervisory controllers (S-PBC) identify phasor targets (relative to a reference) at specific nodes corresponding to a desired power flow solution, where each S-PBC controller is responsible for a zone or portion of the network. The supervisory controller communicates these targets to one or more local controllers (L-PBC) within its zone. L-PBC then recruits one or more local resources to meet the phasor target assigned by S-PBC. Specifically, local controllers communicate suitable commands for actuators to continually modify real and reactive power injections so as to track a phasor target.

2.1. Radial Distribution Feeder

2.1.1. Toy Example 1

A motivating example illustrates how the voltage phasor captures locational aspects of power flow along a distribution circuit. In Figure 1, Resource 1 under PBC injects net real and reactive power P_1 and Q_1 to maintain a target phasor $V_1 \angle \delta_1$ relative to the substation. Any changes at other nodes will affect net power flow on the main feeder, thus instantaneously changing $V_1 \angle \delta_1$. By measuring only two quantities, its local phasor and $V_0 \angle \delta_0$ at the substation, Resource 1 can modulate its output to compensate for these changes. Note that the exact placement of the sensor around the node is not important for measuring voltage.

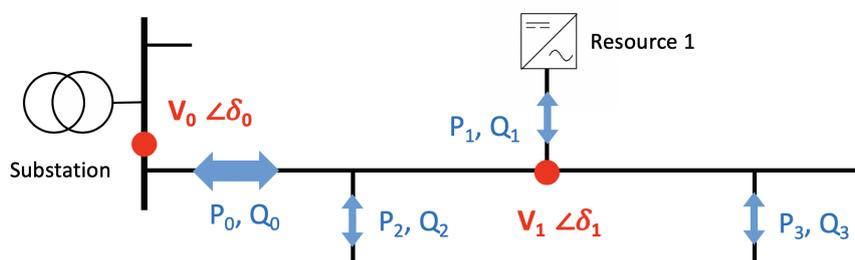


Figure 1. Simple phasor-based control (PBC) example on a typical radial distribution circuit. Two phasor measurement units (μ PMUs) record the voltage phasors $V_0 \angle \delta_0$ and $V_1 \angle \delta_1$ at their respective nodes. Arrows depict real and reactive power flows P and Q , which may be in either direction. Resource 1 represents a controllable power source or sink, such as solar PV, battery, or controllable load, that modulates real and reactive power flows P_1 and Q_1 .

Arguably, one could simply use the net power flow P_0 and Q_0 at the feeder head, which also reflects the sum of changes on the feeder—but this sacrifices the locational information. Consider the case where the net load P_2 drops as the net load P_3 increases, leaving P_0 and Q_0 constant. In this case, line losses and the voltage drop along the feeder increase, potentially causing a violation. Such a violation could be mitigated by restoring the phasor V_1 relative to V_0 , by increasing the P_1 injection that will provide more support near the end of the feeder.

Non-phasor measurements are less useful for managing real and reactive power flow at the same time as voltage constraints. For example, if controlled on a volt-VAR curve, Resource 1 would increase Q_1 in response to declining voltage magnitude V_1 but neglect possible adjustments to P_1 . This can result in significant I^2R losses if real and reactive power flow in opposite directions. Additional current measurements along the feeder could determine the desired P_1 , but a minimum of three additional sensors would be needed. Here, the sensors would measure the voltage magnitude at Resource 1,

the current in the direction of P_0 , the current in the direction of P_2 , and the current in the direction of P_3 . Furthermore, these branch currents are likely to contain more information about individual customers, whereas nodal voltages can be published with less concern about privacy. This enables PBC resources to respond directly to compensate for their neighbors' behavior, without divulging exactly what these neighbors are doing.

Note that a phasor target must be expressed as a difference between two locations, since grid frequency varies with time—that is, phase angle measurements relative to the PMU clock at a single location have no physical meaning. Real-time data transfer is thus required from a reference node in addition to a PMU co-located with the L-PBC controller at the control node. Nevertheless, the communication overhead in this feedback loop is limited to a single non-local channel. Since voltage phasors are associated with network nodes, whereas current and power flow are associated with network branches, a single nodal PMU measurement (e.g., at the substation) becomes a valid reference point for multiple L-PBC controllers along the feeder, or on neighboring feeders, respectively. This architecture further enables extension to nested layering, where a supervisory node for a local aggregation of resources can become a local node for the next higher level of supervisory control, on a larger geographic scale.

2.1.2. Radial Distribution Feeder Simulation

Figure 2 illustrates a simulation detailed in [28], which implements the concept from the preceding toy example. Controllable resources are placed at multiple locations (634, 645, 675) and actuated by L-PBC so as to track the voltage phasor target assigned to the performance node (632) on the Institute of Electrical and Electronics Engineers (IEEE) 13-node test feeder. The phasor target is defined here relative to the reference node (650) across the substation transformer, thus capturing net real and reactive power flow at the feeder head.

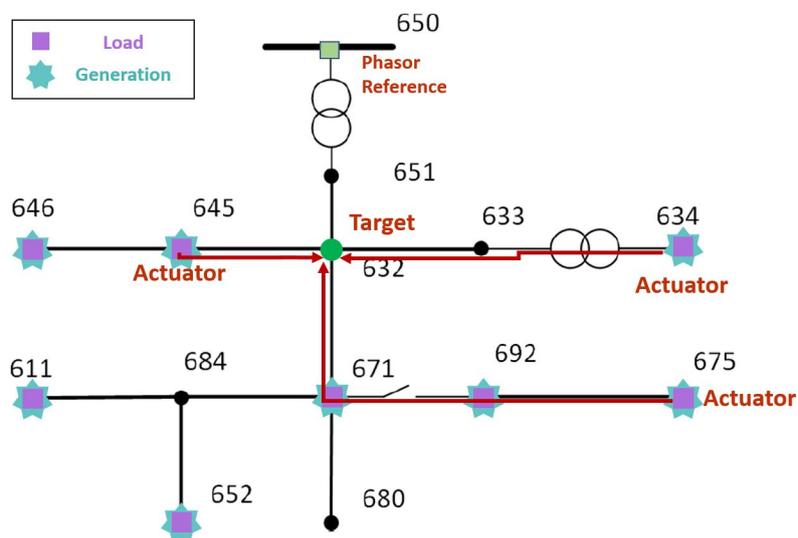


Figure 2. Institute of Electrical and Electronics Engineers (IEEE) 13-node Test feeder [29] for PBC simulation with multiple actuators. PMUs at nodes 650 and 632 measure the respective voltage phasors. The local (L)-PBC controller coordinates actuators at nodes 645, 634, and 675 by modulating their power to track a target phasor reference at Node 632 relative to node 650.

Figure 3 illustrates L-PBC disturbance rejection, whereby controllable resources act in concert. Here, L-PBC is implemented with proportional-integral (PI) control, using two different tuning algorithms. Square wave disturbances representing sudden net load changes elsewhere on the feeder are introduced at 120, 180, and 210 s. For example, the disturbance set on Phase A at 120 s entails a sudden net load increase at node 634 (300 kW + 430 kVar on a 5000/3 kVA per-phase base), 671 (0 kW

+ 200 kVAR), 675 (0 kW + 100 kVAR), and 632 (500 kW + 430 kVAR). Voltage magnitude and phase angle at the performance node are smoothly restored by PBC [28].

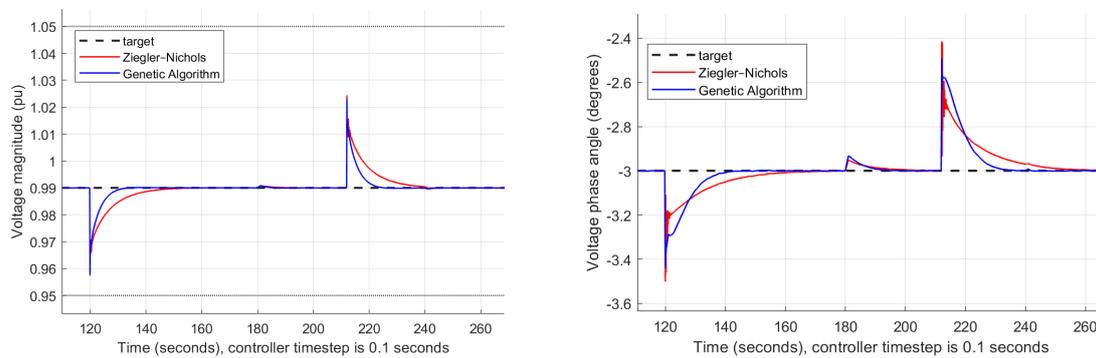


Figure 3. The target phasor at node 634a and the corresponding L-PBC voltage magnitude response (left) and phase angle response (right) to three disturbances. Different algorithms can be used for tuning proportional-integral (PI) controller gains.

2.2. Parallel Transmission Lines

2.2.1. Toy Example 2

Another example motivates the use of phasor targets for stabilizing the network in the event of topology changes. Consider a transmission path with two identical parallel lines connecting Nodes 1 and 2. Suppose one line is suddenly lost, approximately doubling the load on the remaining line (with power demand and supply on either side initially unchanged). The increased line loading may threaten to violate operating constraints, such as thermal ratings, voltage, or angle stability limits, thus mandating some corrective action. Under conventional control, resources at either end would be instructed to change their power injection based on some combination of power flow analysis and economic optimization, based on the new network topology. This takes time to compute and requires information about the nature of the contingency event.

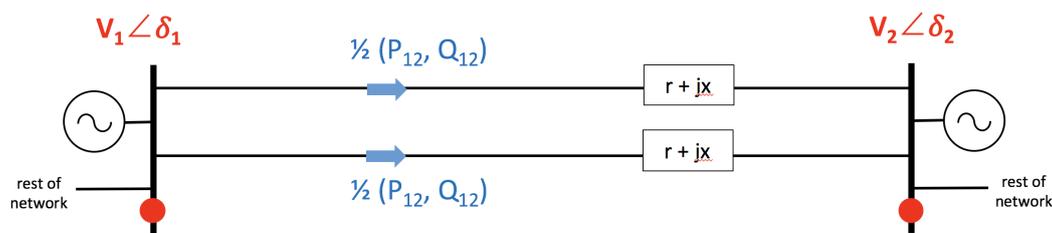


Figure 4. Simple contingency example for a meshed transmission network, where a PMU at each of the two nodes measures the respective voltage phasor. Two transmission lines share a path to deliver the combined power P_{12} , Q_{12} from node 1 to node 2. The impedance of each individual line is $r + jx$. A credible contingency event to be addressed by PBC is the loss of one transmission line.

By contrast, PBC would be tasked with maintaining a scheduled phasor difference $V_1 - V_2$ that is consistent with scheduled power transfers and network constraints. Upon the line trip, the phasor difference instantly jumps, reflecting the doubled impedance between Nodes 1 and 2. Without any need for computation, or information about the contingency event, L-PBC can respond immediately to restore the phasor. For example, if there is excessive flow from node 1 to 2, L-PBC would act to decrease power injection at node 1 and increase generation or decrease load at node 2. (Resources may also be recruited from the “rest of network,” but this is tangential to the argument.)

If sufficient resource actuation is available, the phasor difference will be restored to the original value $V_1 - V_2$, which (by Ohm’s Law) restores the current on the single remaining line to the same

value as before the trip event. The resulting operating state is likely sub-optimal (in terms of the required power injections and their respective costs) but known to be safe with respect to network constraints. Even if the available resources are insufficient for L-PBC to fully restore the phasor in the post-contingency network, the actuation is still guaranteed to drive the phasor in the right direction, toward alleviating the overload and stabilizing the system. This property of PBC is illustrated in more detail below. L-PBC thus buys time for S-PBC to re-optimize the dispatch for the new topology and assign a new set of target phasors.

2.2.2. Simulation: Transmission P-V curves

We generate a P-V curve, which describes the relationship between real power injection and voltage magnitude, for Node 2 in Figure 4. For illustrative purposes, we choose $X = \frac{x}{2}$ and $r \approx 0$, which is typical for transmission systems. Using the full derivation in [30], this transmission-level P-V curve can be expressed as

$$V_2^4 + (2XP_2\beta_2 - V_1^2)V_2^2 + X^2P_2^2(1 + \beta_2^2) = 0, \quad (7)$$

where $\beta_2 = \tan \theta_2 = \frac{Q_2}{P_2}$, and the power factor at node 2 is $\cos \theta_2$.

Figure 5 illustrates the change in this P-V curve under the contingency from Motivating Example 2, with grey lines $\pm 5\%$ about the nominal 120 kV. The sudden loss of one transmission line would reduce the power transfer capability of the path and shrink the P-V curve. With power injection at Node 2 unchanged, this would move the system from operating Point A (1000 kW) to Point B (magnified for clarity in the right side figure), substantially reducing the stability margin. Under PBC, the resource at Node 2 can instantly respond to the contingency by reducing real power demand (Point C) to maintain voltage on the shrunken curve, or by adjusting the power factor, which would shift the P-V curve upward (corresponding to the blue dotted line) and similarly support voltage stability. While the new operating state may be uneconomical to sustain for long periods, it enables the resource to act quickly to help prevent cascading failures.

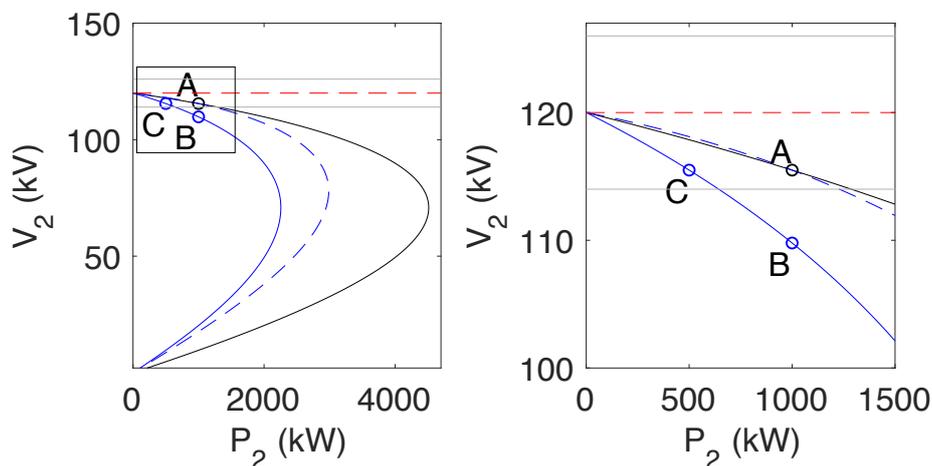


Figure 5. Power-Voltage stability curves for a contingency event (magnified for clarity in the right side figure). Operating point A is pre-contingency; points B and C are post-contingency. Without actuation, the loss of a parallel transmission line shifts the curve and forces the system from A to B. By explicitly addressing nodal voltage, PBC drives the system toward operating point C. This avoids dangerous proximity to the nose of the curve, allowing for an improved stability margin without requiring calculation of the new curve.

3. The PBC Paradigm: Control System Architecture

3.1. Local PBC

For local phasor tracking, each L-PBC controller drives one or more DER to meet or restore the S-PBC assigned phasor target, with feedback provided by phasor measurements. Along with the targets communicated by the supervisory layer, the local controller inputs are direct empirical measurements of voltage magnitude V and angle $\angle\delta$ at a local performance node and a reference node, and its outputs are real and reactive power (P and Q) commands to actuators located at or near the performance node.

Our extended project team has developed and tested multiple L-PBC algorithms, including proportional-integral (PI) control [28], Linear Quadratic Phasor Control (LQPC), extremum-seeking (ES) control [31], and Retrospective Cost Adaptive Control (RCAC) [32]. The details of the controller design are not central to the PBC paradigm. The various approaches are summarized here to establish the general feasibility of PBC, given challenges such as the nonlinearity of the power flow equations and limited information about the network model.

The simple PI controller has the advantage of being transparent and tunable. The design objective is to determine a robust strategy for tuning the controller gains, either on-line or off-line, to ensure proper actuation that achieves relatively rapid convergence on the phasor target, while preventing oscillations or adverse interactions among multiple local controllers. Since P and Q are related to both V and δ , the tuning algorithm uses grid characteristics to consider these cross interactions when implementing each respective controller feedback loop [28].

The LQPC, a multiple-input/multiple-output (MIMO) controller, suits the PBC problem more elegantly. The LQPC is a modified version of a Linear Quadratic Regulator (LQR), which defines its cost function as a linear combination of a quadratic phasor error term and a quadratic actuator (P and Q injection) effort term. The quadratic weight matrices in the cost function provide intuitive tuning knobs for the LQPC controller. Using the cost functions and a linearized power flow model for the network, the LQPC uses the LQR Riccati equations to provide an infinite-horizon linear time-invariant MIMO feedback matrix for both the P and Q injections, based on the measured phasor error. To compensate for error in the linearized power flow model, the LQPC incorporates error integration states for both V and δ that eliminates steady-state error.

The RCAC controller, presented in [32], is the newest of the algorithms under study for this application. The RCAC controller uses an impulse response as an empirical method for modeling the relationship between P, Q injections at the actuation node and the voltage phasor at the performance node. This controller does not rely on any externally provided information about impedances within the network.

3.2. Supervisory PBC

The S-PBC controller performs an optimization—in essence, an optimal power flow (OPF) calculation. Input information for S-PBC includes technical and cost data for participating (controlled) resources, forecasts for loads and other non-controlled resources, information about the network (including topology, impedances, constraints, and real-time operating state), and a statement of the objective function that reflects the desired optimization criteria. Some of the S-PBC inputs, including system state and network impedance parameters, may be estimated from PMU measurements [33]. The output is an expression of the desired steady state power flow for the entire network or a region thereof, in terms of target phasors at each node. Since nodal voltage phasors are intrinsic state variables for any a.c. power flow calculation, the main novelty here lies in reporting explicitly the phasor values that would otherwise remain “under the hood” of the OPF formulation.

Our PBC scheme allows supervisory controllers to operate on different timescales as needed. On each scale, the S-PBC controller determines the collection of target voltage phasors for the local control layer by solving a constrained optimization-based problem, but the geographic scope and

comprehensiveness of the optimization can be flexible. When target phasors need to be updated in response to rapidly evolving local network conditions, we envision that S-PBC action, followed by L-PBC action and device response, could occur within seconds. The operational adjustments on such a fast time step would be small increments, consistent with constraints on resource ramp rates that are enforced by each local controller. More typically, the S-PBC controller would update phasor targets on the order of minutes or even hours, as warranted by significant changes in operating conditions or costs that call for comprehensive re-optimization. The slower timescale integrates transmission and distribution tier objectives with power flow solutions to improve the overall performance of PBC. Thus, S-PBC does not require updating the full supervisory optimization with every phasor target update, but it will always respect local constraints.

The essential property of this layered strategy is that the desired operating state of the system is continually approximated by prioritizing harder and more local constraints before refining the optimization with additional, more global information. Aside from the lower bound on the latency and response capability of local controllers and actuators, nothing in the PBC scheme dictates a particular choice of time step for optimization; nor do multiple controllers need to operate on the same step.

In principle, the structure of the hierarchical controller admits any method for obtaining an optimal power flow solution at the supervisory level. One appealing choice is the three-phase linear unbalanced power flow model developed in [24]. This linear approximation can be used as a set of constraints that enforce the physics of the grid in a quadratic program, which is quickly solved even for large network models.

4. Scalable Integration of Variable Energy Resources

4.1. Managing Variability

When a considerable proportion of electricity is supplied from renewable energy sources, such as wind turbines and solar PV, we expect higher variability of power flows and distribution supply voltages—seasonally, daily, and minute-to-minute (e.g., due to fast moving clouds or varying wind speeds). By setting explicit phasor bounds, PBC ensures the limits of the infrastructure are respected, even as power and voltage quantities vary within the permissible range. In other words, the emphasis on non-negotiable physical constraints allows PBC to flexibly accommodate evolving solutions for how the constraints are met.

The precision with which target phasors are maintained, or the deadband settings for local controllers, can be adapted to the requirements of each particular situation. A key feature of the PBC framework is that it allows for feedback at both the local and supervisory level. If, for any reason, an L-PBC controller is unable to recruit the resources needed to meet its phasor target with the desired accuracy, the discrepancy—in the form of a direct phasor measurement—is directly apparent to the S-PBC controller, which can then revise phasor targets to best manage the actual condition. PBC thus accounts for variability by design, from rejecting disturbances in real-time to adapting to changes in resource availability. We believe this makes the PBC framework ideally suited to address the fluctuations and uncertainty associated with high levels of variable energy resources.

4.2. Layered Architecture

The PBC framework, per se, is agnostic to the optimization objective, criteria, or method applied when computing phasor targets, as well as the specific means employed by local controllers to meet their targets. For example, S-PBC may run a single optimal power flow program or allow for some market-based process in determining target phasors. L-PBC enforces S-PBC target voltage phasors by measuring the actual phasor in real-time and modulating power to and from controllable energy resources—tracking the assigned target phasor. L-PBC may itself optimize locally based on information about specific resources (such as marginal cost, battery state-of-charge, etc.) in order to meet its given target phasor. Using nodal phasors as the explicit control variable allows for great

diversity of implementation across controllers and zones because the only requirements for consistency are precisely the physical boundary conditions that are spelled out as phasors.

The PBC framework is also agnostic about the control time step, presuming only that L-PBC follows phasor targets faster than S-PBC updates them. At the local level, the most immediate goal is to compensate for disturbances to ensure a safe and secure operating state. Local control can be fast because it requires very limited information: a local phasor measurement, a reference phasor measurement, and any change in status among the resources under its control. The phasor measurements will instantaneously reflect changes in system conditions, which constitute disturbances to be rejected by the local resources. For example, resources under PBC might be actuated once per second to compensate for changes in load or the behavior of uncontrolled resources. (Our local controller for the simulation in [28] used a time step of 0.1 s.) By acting to restore the phasor, the local controller will drive the system toward an operating state known to be safe and stable, without information about the nature of the disturbance. Fundamentally, this is because both power flows and physical operating constraints of the a.c. grid are directly expressible in terms of phasor differences between nodes.

With the luxury of additional information and time, supervisory control algorithms can always update phasor targets to account for additional criteria as desired or realistic under the given circumstances in order to optimize the system. The temporal disaggregation—maintaining a technically feasible operating state first, and then adjusting targets to fine-tune the system—inherently prioritizes grid stability and reliability. This agrees with the intent of traditional grid operations but reverses the standard approach: today, optimization comes first, followed by a check for possible violations that may overrule the desired solution. By starting from explicit physical constraints, PBC promises to be more robust, while avoiding costly over-corrections.

Stacked PBC layers not only offer a scalable path to implementation but also a way to integrate performance goals between distribution and transmission tiers. For example, the objective of a supervisory controller at a distribution substation might be to balance three-phase voltages while minimizing the volatility that creates unnecessary wear on legacy voltage regulation equipment. Accordingly, the S-PBC controller computes and assigns target phasors to strategically located L-PBC controllers along the feeder. The L-PBC controllers then coordinate distributed energy resources, such as loads, generators, or storage, to achieve the S-PBC objective, while rejecting disturbances on a much faster time-scale. The S-PBC controller may then additionally negotiate with a transmission-level entity to provide frequency regulation or other ancillary services. S-PBC would then re-compute and assign L-PBC phasor targets to meet the additional criteria, with the grid performance from the transmission tier incorporated. Other criteria, such as individual resource cost functions, feeder loss minimization, etc., are straightforward to include by design.

We propose that the formulation of optimal power flow objectives in terms of nodal phasors is generalizable across transmission networks. Just as large generators today modulate reactive power output to maintain a target bus voltage magnitude, we envision generators adjusting both real and reactive power output to track a complete bus voltage phasor. This logical extension was historically inconceivable without a way to empirically measure the timing of the voltage waveform (where the phase angle translates to the rotor position) with sufficient accuracy. Today, high-precision phasor measurements can serve as a feedback control input for both reactive and real power, which could effectively replace velocity with position control for rotating machines [34].

4.3. Strategic Advantages of PBC

By assigning target phasors in terms of the grid state variables V and δ , the supervisory layer manages power flows at any chosen node relative to another on the grid, without articulating how to attain that prescribed condition (e.g., through a particular dispatch of resources). We believe that this approach has unique strategic value beyond the integration of variable renewable resources,

particularly in helping prevent cascading failures and supporting improved restoration processes for grid resilience.

As illustrated by the above examples, a scheduled phasor profile will remain a feasible operating state in the event of contingencies—including network topology changes—whereas a scheduled power dispatch may not. This is because the phasor-based formulation of grid operating objectives focuses on the essence of what is physically required to keep the grid stable, while being agnostic to the logistics of how these objectives are met. Crucially, the use of the phasor as a network state variable enables an immediate corrective response without communication among controllers.

Suppose that all phasor targets prescribed by S-PBC are met, and the system is operating optimally. Any subsequent disturbance or change in load will render these targets suboptimal, but they will still meet system constraints and represent safe operation until new targets are computed. That is, the L-PBC controller mitigates load and generation variability by tracking the assigned S-PBC phasor. When an L-PBC controller fails to meet its target, this condition will be transparent to other controllers insofar as they will observe a change in the voltage phasor they are monitoring. Such an event can trigger the S-PBC controller to recompute the target phasor, accounting for increased variability in the network segment. To distinguish between cases where the L-PBC controller fails to meet a target as a result of (a) increased load and generation variability; (b) actuator saturation (e.g., the battery is fully charged and cannot absorb more energy); or (c) actuation failure (e.g., the PV inverter malfunctions), we may rely on (1) updated forecasts of load and generation; (2) the status of the integrator state and employ an integrator anti-wind-up scheme; and/or (3) a status report from the actuator.

The discussion on voltage stability above illustrates that control actions toward restoring assigned phasors after disturbances and contingency events will drive the system state in the direction of stability, or away from the “nose” of the curve, on $P-\delta$, $P-V$, or $Q-V$ curves for any node. Of course, conventional operating decisions would hope to drive the system in the same direction. The crucial difference is that the “right direction” is not obvious when the operating state of the grid is primarily expressed in terms of power injections and flows, and the state variables are computed or measured as an afterthought. Thus, in conventional practice, information about which control action at any given node would drive the system toward or away from stability is not readily available, especially to local resources. By contrast, two pieces of information—the phasor target assigned by S-PBC, and a pair of real-time phasor measurements—are sufficient to empower L-PBC controllers to “do the right thing” at their respective nodes. Actuation toward restoring phasors will further help prevent cascading sequences, where changes to P or Q inadvertently exacerbate a disturbance to the system state, as observed at a neighboring node.

Another strategic advantage of the phasor-based formulation is that it can support safe and speedy restoration processes, including “self-healing” functionality. Specifically, aligning phasor targets for interconnections or between different parts of the network enables the opening and closing of switches under load. When the main concern is the amount of arcing between electrical contacts, limiting the phasor difference between two sides of an open switch is a direct way to ascertain safe operation. The same rationale applies to high-voltage transmission switchgear, as well as sectionalizing switches in distribution systems or points of common coupling for microgrids.

The use of explicit phasors minimizes the information required for intelligent local operational decisions. For example, if the objective is to drive the phasor difference across a tie switch to within a safe value for closing that switch, the supervisory controller can simply assign the appropriate phasors on either side (or, if only one side is controllable, assign it to match the uncontrolled side) and leave the details of resource actuation to the L-PBC layer. By focusing on the most essential information about the grid operating state, PBC obviates the need for more centralized and cumbersome communication and data management structures.

Nevertheless, the possibility of communication loss must be accounted for by design. Even though a local controller requires only a single non-local piece of information (the reference phasor) to hold a target, it must have a default fail-safe mode in case the communication channel drops out.

Possible strategies include minimizing departures from the most recent operating state known to be compliant and resorting to managing only voltage magnitude (which is strictly local). In addition, PBC is admittedly vulnerable to loss or spoofing of the GPS pulse-per-second signal. Because of the growing importance of PMUs in power system operations worldwide, alternative technologies to GPS timekeeping are already being considered and could plausibly meet the future needs of a control infrastructure that will depend more heavily on secure time synchronization. [35]

Arguably, the PBC paradigm could improve grid resilience with respect to cyber-attacks by virtue of its inherent simplicity and transparency. The crucial information to be transacted by controllers and participating resources—namely, the physical operating state of the grid—is both minimal in quantity and independently verifiable by any interested party. For example, phasor differences across the transmission and distribution system are easily observed with μ PMUs from any wall outlet, whereas current and power flows are measurable only on their particular circuit branch. While malicious cyber-intrusions can never be ruled out, the widespread use of explicit phasor measurements should make tampering with operationally relevant data more readily detectable. We look forward to further study of cyber-security implications and optimally secure design of PBC.

5. Conclusions

The PBC framework is a network-centered rather than resource-centered control paradigm that uses the state variables—voltage magnitudes and phase angles at specific network nodes—as explicit control variables. This approach aims to maintain a steady-state power flow condition with local disturbance rejection, giving distributed resources the information necessary to act quickly and semi-autonomously in the interest of stabilizing the network. PBC is fundamentally agnostic to the control time step, the nature of the resources recruited, or the optimization criteria for determining the desired steady-state. The PBC framework simultaneously addresses multiple operational challenges, including high resource variability, reverse power flow, grid visibility, and coordination between transmission and distribution tiers. We anticipate that PBC can facilitate the deployment of high levels of variable and distributed power generation, while supporting grid reliability and resilience. In particular, PBC would solve the problem of voltage management for distribution utilities when the only recourse otherwise is to limit feeder hosting capacity for DER. At the same time, by supporting coordinated real and reactive power balance from the periphery of the grid, PBC will mitigate the need for utilities and system operators to purchase expensive reserves to assure frequency and voltage stability. Relying on such reserves may become more difficult, costly, and less reliable as renewable penetration increases.

Our team's initial simulation results suggest that PBC will be feasible to carry out at the distribution level and scalable across the transmission level. Further simulation and hardware-in-the-loop testing is ongoing to address specific design challenges and to build a solid empirical basis for assessing the potential practical impact of this radically new control paradigm.

Author Contributions: Conceptualization A.v.M.; Analysis E.R., K.B., J.S. and K.M.; funding acquisition A.v.M. and E.R.; investigation K.B., J.S. and K.M.; methodology A.v.M. and E.R.; project administration A.v.M. and E.R.; supervision A.v.M. and E.R.; validation K.B., J.S. and K.M.; visualization E.R. and J.S.; writing (original draft) A.v.M. and E.R.; writing (review & editing) J.S. and K.M. All authors have read and agreed to the published version of the manuscript.

Acknowledgments: The authors wish to thank their colleagues, including Daniel Arnold, Christoph Gehbauer, Michael Sankur and Evangelos Vrettos (Lawrence Berkeley National Laboratory), Dennis Bernstein, Adam Bruce and Syed Aseem Ul Islam (University of Michigan), Ravi Venugopal, Sudipta Chakraborty and Guna Bharati (OPAL-RT), Ali Vojdani, Terry Nielsen and Shangyou Hau (GridBright), Harby Sehmar and John Mead (Pacific Gas & Electric Co.), Sean Patrick Murphy (PingThings), Kevin D. Jones (Gridwidget Labs and Dominion Energy), Moustafa Abdelbaky, Leo Chu, Gabe Fierro, Jasper Pakshong, Jiasheng Qin and T.G. Roberts (University of California, Berkeley) and Carl Blumstein and Eric Lee (California Institute for Energy and Environment) for their important contributions to creating a successful research project. We also thank Dave Auslander (U.C. Berkeley) for helpful conversations.

Funding: This research was funded by U.S. Department of Energy, Award DE-EE0008008.

Conflicts of Interest: The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

References

1. Markard, J. The next phase of the energy transition and its implications for research and policy. *Nat. Energy* **2018**, *3*, 628–633.
2. Mead, J.; Donde, V.; Garnett, J. Advanced Control Technologies for Distribution Grid Voltage and Stability With Electric Vehicles And Distributed Generation. Available online: <https://www.energy.ca.gov> (accessed on 11 May 2019).
3. Australian Energy Market Operator Limited. Technical Integration of Distributed Energy Resources: Improving DER Capabilities to Benefit Consumers and the Power System. A Report And Consultation Paper (2019). Available online: <https://www.aemo.com.au> (accessed on 11 May 2019).
4. Ratnam, E.L. Balancing Distributor and Customer Benefits of Battery Storage Co-Located With Solar PV. Available online: <http://hdl.handle.net/1959.13/1312195> (accessed on 11 May 2019).
5. Xu, Y.; Çolak, S.; Kara, E.C.; Moura, S.J.; González, M.C. Planning for electric vehicle needs by coupling charging profiles with urban mobility. *Nat. Energy* **2018**, *3*, 484–493.
6. Fares, R.L.; Webber, M.E. The impacts of storing solar energy in the home to reduce reliance on the utility. *Nat. Energy* **2017**, *2*, 17001 .
7. Bogdanov, D.; Farfan, J.; Sadovskaia, K.; Aghahosseini, A.; Child, M.; Gulagi, A.; Solomon, A.O.; de Souza Noel Simas Barbosa, L.; Breyer, C. Radical transformation pathway towards sustainable electricity via evolutionary steps. *Nat. Commun.* **2019**, *10*, 1077 .
8. Wenzel, G.; Negrete-Pincetic, M.; Olivares, D.E.; MacDonald, J.; Callaway, D.S. Real-Time Charging Strategies for an Electric Vehicle Aggregator to Provide Ancillary Services. *IEEE Trans. Smart Grid* **2018**, *9*, 5141–5151.
9. Davies, D.M.; Verde, M.G.; Mnyshenko, O.; Chen, Y.R.; Rajeev, R.; Meng, Y.S.; Elliott, G. Combined economic and technological evaluation of battery energy storage for grid applications. *Nat. Energy* **2019**, *4*, 42–50.
10. Kittner, N.; Lill, F.; Kammen, D.M. Energy storage deployment and innovation for the clean energy transition. *Nat. Energy* **2017**, *2*, 17125.
11. Chen, X.; Zhang, H.; Xu, Z.; Nielsen, C.P.; McElroy, M.B.; Lv, J. Impacts of fleet types and charging modes for electric vehicles on emissions under different penetrations of wind power. *Nat. Energy* **2018**, *3*, 413–421.
12. von Meier, A.; Stewart, E.; McEachern, A.; Andersen, M.; Mehrmanesh, L. Precision Micro-Synchrophasors for Distribution Systems: A Summary of Applications. *IEEE Trans. Smart Grid* **2017**, *8*, 2926–2936.
13. von Meier, A.; Culler, D.; McEachern, A.; Arghandeh, R. Micro-synchrophasors for distribution systems. In Proceedings of the IEEE PES Conference on Innovative Smart Grid Technologies, Washington, DC, USA, 19–22 February 2014; pp. 1–5.
14. Lund, P.D.; Lindgren, J.; Mikkola, J.; Salpakari, J. Review of energy system flexibility measures to enable high levels of variable renewable electricity. *Renew. Sustain. Energy. Rev.* **2015**, *45*, 785–807.
15. Deng, R.; Yang, Z.; Hou, F.; Chow, M.; Chen, J. Distributed Real-Time Demand Response in Multiseller–Multibuyer Smart Distribution Grid. *IEEE Trans. Power Syst.* **2015**, *30*, 2364–2374.
16. Álvaro Ortega.; Milano, F. Frequency Control of Distributed Energy Resources in Distribution Networks. *IFAC-PapersOnLine* **2018**. *51*, 37–42.
17. Cheng, B.; Powell, W.B. Co-Optimizing Battery Storage for the Frequency Regulation and Energy Arbitrage Using Multi-Scale Dynamic Programming. *IEEE Trans. Smart Grid* **2018**, *9*, 1997–2005.
18. Nourai, A.; Schafer, C. Changing the electricity game. *IEEE Power Energy Mag.* **2009**, *7*, 42–47.
19. Li, R.; Wu, Q.; Oren, S.S. Distribution Locational Marginal Pricing for Optimal Electric Vehicle Charging Management. *IEEE Trans. Power Syst.* **2014**, *29*, 203–211.
20. Smith, J.W.; Sunderman, W.; Dugan, R.; Seal, B. Smart inverter volt/var control functions for high penetration of PV on distribution systems. In Proceedings of the 2011 IEEE/PES Power Systems Conference and Exposition, Phoenix, AZ, USA, 20–23 March 2011; pp. 1–6.
21. Kroposki, B.; Johnson, B.; Zhang, Y.; Gevorgian, V.; Denholm, P.; Hodge, B.; Hannegan, B. Achieving a 100% Renewable Grid: Operating Electric Power Systems with Extremely High Levels of Variable Renewable Energy. *IEEE Power Energy Mag.* **2017**, *15*, 61–73.

22. Ochoa, L.; Wilson, D. Angle Constraint Active Management of Distribution Networks with Wind Power. In Proceedings of the 2010 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT Europe), Gothenberg, Sweden, 11–13 October 2010.
23. Eftekharijad, S.; Vittal, V.; Heydt, G.T.; Keel, B.; Loehr, J. Impact of increased penetration of photovoltaic generation on power systems. *IEEE Trans. Power Syst.* **2013**, *28*, 893–901.
24. Sankur, M.D.; Dobbe, R.; Stewart, E.; Callaway, D.S.; Arnold, D.B. A Linearized Power Flow Model for Optimization in Unbalanced Distribution Systems. *arXiv* **2016**, arXiv:1606.04492.
25. Baran, M.E.; Wu, F.F. Network reconfiguration in distribution systems for loss reduction and load balancing. *IEEE Trans. Power Del.* **1989**, *4*, 1401–1407.
26. Dobbe, R.; Arnold, D.; Liu, S.; Callaway, D.; Tomlin, C. Real-Time Distribution Grid State Estimation with Limited Sensors and Load Forecasting. In Proceedings of the 2016 ACM/IEEE 7th International Conference on Cyber-Physical Systems (ICCPS), Vienna, Austria, 11–14 April 2016; pp. 1–10.
27. Moffat, K.; Bariya, M.; von Meier, A. Delta Linearized Power Flow. Power Systems Controls Conference (PSCC) (in review), 2020.
28. Swartz, J.; Roberts, T.; von Meier, A.; Ratnam, E.L. Local Phasor-Based Control of DER Inverters for Voltage Regulation on Distribution Feeders. In Proceedings of the IEEE GreenTech Conference, Oklahoma City, OK, USA, 1–3 April 2020 (in press); pp. 1–7.
29. IEEE Distribution Test Feeders. 13-Bus Feeder. Available online: <http://sites.ieee.org/pes-testfeeders/resources/> (accessed on 30 April 2018).
30. Bhaladhare, S.B.; Telang, A.S.; Bedekar, P.P. P-V, Q-V Curve—A Novel Approach for Voltage Stability Analysis. *IJCA Proc. Natl. Conf. Innov. Paradig. Eng. Technol.* **2013**, *5*, 31–35.
31. Sankur, M.; Dobbe, R.; von Meier, A.; Arnold, D. Model-Free Optimal Voltage Phasor Regulation in Unbalanced Distribution Systems. *IEEE Trans. Smart Grid* **2019**, *11*, 884–894.
32. Islam, S.A.U.; Ratnam, E.L.; Goel, A.; Bernstein, D.S. Phasor-Based Adaptive Control of a Test-Feeder Distribution Network: Application of Retrospective Cost Adaptive Control to the IEEE 13-Node Test Feeder. *IEEE Control Syst. Mag.* **2019**, *39*, 56–74.
33. Moffat, K.; Bariya, M.; von Meier, A. Real Time Effective Impedance Estimation for Power System State Estimation. In Proceedings of the IEEE PES Conference on Innovative Smart Grid Technologies (accepted), 2020.
34. Auslander, D. (University of California, Berkeley). Phase-Controlled Synchronous Generators for a Constant Frequency Grid. Manuscript in preparation, 2019.
35. Kelly, C.; Pellegrino, J.; Taylor, E.; Goldstein, A. *Time Distribution Alternatives for the Smart Grid*. Technical Report; National Institute of Standards and Technology (NIST): Gaithersburg, MD, USA, 2017.



© 2020 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).