Phasor-Based Control for Scalable Solar PV Integration

Final Project Report
Review Draft January 27, 2021

Prepared for
U.S. Department of Energy's Office of
Energy Efficiency and Renewable Energy (EERE)
Solar Energy Technologies Office (SETO)

Prepared by
California Institute for Energy and Environment
(CIEE)

In coordination with
University of California, Berkeley
Lawrence Berkeley National Laboratory
University of Michigan
OPAL-RT Corporation
GridBright, Inc.
PingThings, Inc.

2150 Allston Way, Suite 280
Berkeley, CA 94704-110
(510) 664-7462 | uc-ciee.org

Acknowledgement:

This material is based upon work supported by the Department of Energy, Office of Energy Efficiency and Renewable Energy, Solar Energy Technologies Office, under Award Number DE-EE0008008.

Disclaimer:

Executive Summary

This project demonstrated that solar PV can be recruited to stabilize the grid, smooth out disturbances, manage power flows, and assist circuit switching operations. It developed a radically new, layered control framework for Distributed Energy Resources (DER) to act in response to real-time, measured conditions on their local distribution circuit, rather than waiting for a price signal to indicate preferred behavior. By enabling resources to act as good citizens on the electric grid, Phasor-Based Control (PBC) facilitates arbitrarily high solar penetration levels.

PBC expresses objectives in terms of *voltage phasors,* which include information about both the magnitude and the precise timing of grid voltage at each specific location. A supervisory (S-PBC) controller sets voltage phasor targets at different nodes in the transmission or distribution network, and local (L-PBC) controllers recruit real and reactive power from resources such as solar inverters, batteries or loads to track phasor targets. The technology makes use of ultra-precise measurements from micro-phasor measurement units (μPMUs).

Phasor-based control prioritizes stabilizing the grid locally, toward operating states known to be safe in accordance with physical operating constraints, while buying time for economic re-optimization after major changes or contingencies. In doing so, it advances grid reliability and resilience. The framework supports many diverse use cases that specify desired voltage phasors at certain nodes. It applies to distribution as well as transmission systems, although this project focused primarily on distribution applications and simulation. Sample use cases tested in this project include power flow control, voltage management, phase balancing, and support for switching operations.

Key Project Accomplishments include the following:

- Built a repository of non-proprietary test circuit models.
- Demonstrated that supervisory phasor-based control can produce phasor targets consistent with optimal power flow objectives, using suitable linearization techniques.
- Demonstrated that local phasor-based controllers can be effectively tuned to recruit distributed resources for tracking a target phasor and reject disturbances.
- Built a novel communications and control infrastructure, termed the Distributed Extensible Grid Control (DEGC) platform.
- Demonstrated fast simulation capabilities for large networks using a novel partitioning strategy.
- Produced successful PBC simulations on large, realistic distribution circuits.
- Validated simulation results in hardware-in-the-loop (HIL) testing with physical μPMUs, inverters and load banks.
- Produced a value analysis for PBC based on enhanced grid security.
- Identified proliferation opportunities for PBC and infrastructure technology.

The technology is now ready to be field deployed in an experimental pilot setting.

Goals, Objectives and Actual Accomplishments of the Project

This project proposed to develop a radically new layered control framework termed Phasor-Based Control (PBC) for managing extremely high (>100%) penetrations of solar generation and other variable energy resources. PBC provides a framework for Distributed Energy Resources (DER) to act in response to real-time, measured conditions on their local distribution circuit, rather than waiting for a price signal to indicate preferred behavior. This control paradigm prioritizes physical constraints over economic optimization and allows DER to be good citizens on the grid under a wide range of conditions, including contingencies.

Phasor-based control organizes control objectives for an a.c. network by referring explicitly to the state variables V and delta (the voltage phasors) at each node that uniquely determine real and reactive power flows. One or more supervisory controllers (S-PBC) identify phasor targets (relative to a reference) at specific nodes that correspond to a desired power ow 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. Local controllers then recruit one or more local resources to meet the assigned phasor target at their node, by communicating suitable commands to modify real or reactive power injection. Zones may be layered or nested such that control is delegated to supervisory controllers for subordinate zones, subject to phasor continuity at the zone boundaries.

While the most attractive value proposition for PBC is likely found at the bulk transmisison level, the goal of this project was to build PBC from the bottom up. The project scope allowed for detailed algorithm development and hardware-in-the-loop testing at the scale of an individual distribution circuit, along with an initial exploration of scaling in a pure simulation environment, to ascertain that the framework is fundamentally feasible.

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 may itself optimize locally based on information about specific resources (e.g. marginal cost, constraints, etc.) in order to meet its given target phasor, using a suitable feedback controller that measures the actual voltage phasor at the control node relative to the reference. 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 being 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. The L-PBC loop 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 reect 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.

The work integrates several cutting-edge threads of research and development including high-precision voltage phasor measurements, analytics relating phasor profiles to dynamic and unbalanced power flows, decentralized adaptive control algorithms, and simulation capabilities to effectively characterize large networks with heterogeneous, variable and distributed energy resources.

Actual accomplishments over the course of the forty-month project period are compared to the specific objectives from the Statement of Project Objectives (SOPO) within each of the following task groups below:

- Case Design
- Control Algorithm Development
- Controller Implementation and Data Infrastructure
- Hardware in the Loop (HIL) Testing and Simulation

In sum, the project developed the theoretical basis for PBC, created and exercised a range of simulation scenarios and test cases, wrote algorithms for supervisory and local control that meet specific challenges, built a testbed to validate the performance of different PBC algorithms that included physical equipment linked by a highly performant data infrastructure, and succeeded overall in proving the concept: we have established that PBC works. We have also learned what is difficult about it, and how to identify edge cases where it may not work well. The technology is now at a readiness stage where it can be field deployed in an experimental pilot setting.

Case Design

The goal of this task group was to operationalize the idea of PBC in terms of testable hypotheses to determine whether local and supervisory control based on synchrophasor measurements is workable in principle, and whether Phasor-Based Control (PBC) will be capable of achieving the desired objectives to enable 100% solar penetration. For this purpose, we defined a range of PBC scenarios along with performance metrics. Case design included selecting representative distribution circuits, defining constraints, objectives and optimization criteria, and identifying device and data requirements. This effort ensured that the project effectively demonstrates the capability of PBC to meet the key performance targets specified in the FOA, with special emphasis on achieving 100% and greater penetration levels of solar PV generation on distribution circuits.

In consultation with PG&E and other industry advisors, the team identified a set of four objectives for case design that would highlight the capabilities of PBC: namely, power flow control (e.g., to prevent reverse flow on a radial distribution feeder), voltage volatility mitigation (e.g., to reduce disturbances caused by uncontrolled solar PV), phase balancing (e.g., to reduce the impact of single-phase connected PV on power quality and reliability), and phasor alignment or target matching (e.g., to enable flexible circuit switching operations). These were detailed in Milestone 1.1.1 Report. Simulations throughout the project demonstrated that each of these objectives could be achieved with PBC.

The first year efforts focused on the selection of representative circuits, defining constraints, PBC objectives and optimization criteria, building and validating a single feeder, and fully defining the PBC test cases. The objective was to have 30 circuit models identified and available for PBC testing at the end of the first year. In the second year, the case design task aimed to scale

up the modeling and validation of the circuit models in preparation for the OPAL-RT simulation effort targeting 1,000,000 virtual nodes. By the end of the second year, sufficient circuit models were to be built to indicate scalability to 1,000,000 virtual nodes. In the third year, circuit models were built primarily to support the HIL testing and simulation effort at the LBNL FLEXLAB.

The project met its target of 30 distribution circuit models curated, which were made publicly available by GridBright on www.bettergrids.org. The selection included a variety of circuits of different size, complexity, and phase imbalance, ranging from four nodes to several thousand nodes. Thirteen models were adapted from a set of anonymized feeders from PG&E, which required conversion from GridLAB-D to ePHASORsim formats. These feeder models were "realistic" not only in terms of the physical circuits represented, but also in the occurrences of missing detail or inconsistencies (for example, in the representation of phase connectivity or the behavior of voltage regulators) for which real-world distribution circuit models are generally famous. An additional outcome of the project was the enhancement of the DiTTo model conversion tool (originally developed by NREL) to provide better automated conversion and correction functionality. Nevertheless, circuit model conversions still required manual effortson the part of GridBright and OPAL-RT.

Time-varying loads and PV generation profiles were created and superimposed on the circuit models, drawing on a combination of real and synthetic data. These load profiles were selected or created to include large variations or load volatility, with big step changes intended to challenge the PBC controllers, and with PV penetrations of up to 125% of coincident peak load.

Using 50 of these circuits, OPAL-RT produced a transmission and distribution interconnection model in the phasor domain, where the transmission network and distribution circuits are stitched together at the substations using dynamic Thévenin equivalents which allow simulating the propagation of disturbances across the T&D interface. OPAL-RT demonstrated a 185,000-node T&D integration in real-time, and a 660,000-node model offline. While the project did not have the resources or infrastructure to conduct a full integrated PBC simulation on such models, which far exceed the capacity of the OPAL-RT simulator at the FLEXLAB, two important facts were demonstrated: first, that supervisory PBC was capable of performing its optimization at the transmission level, and second, that the methodology scales with reasonable computation times to allow dynamic simulations for very large networks. The results indicated that applying the same architectural rules and adding more parallel computing cores would feasibly permit scaling the simulation methodology to 1,000,000 nodes, the target posited in the FOA.

Control Algorithm Development

The overall PBC strategy is a combination of supervisory control (S-PBC) and local feedback control (L-PBC), where S-PBC addresses the 'Enhanced System Level' in the FOA, and L-PBC the 'Local Device Level' requirement. In this hierarchical control architecture, the supervisory layer provides the target phasor for each node and the local layer modulates the active and reactive power of each controlled device (such as a smart inverter or controllable load) in order to ensure that the reference phasor is accurately tracked at each node. The high level PBC system overview is shown in Fig. 1.

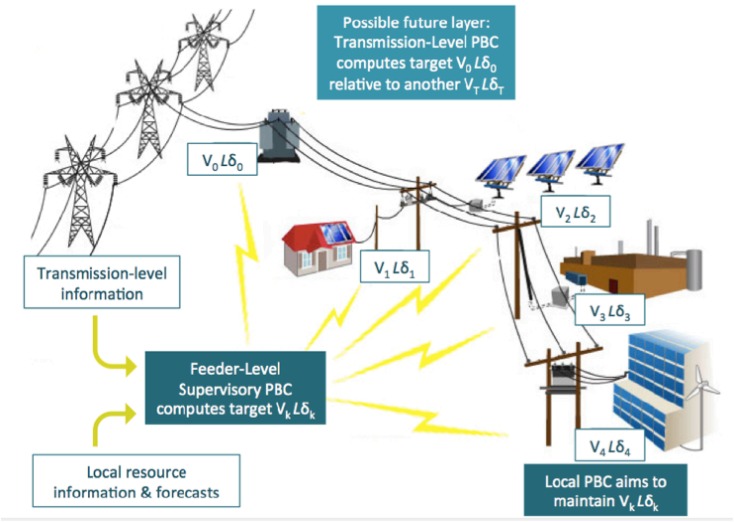

Figure 1: Phasor-Based Control Overview.

The algorithm development tasks comprised two distinct research efforts: namely, the assignment of phasor targets for network optimization (S-PBC), and the generation of real and reactive power commands to individual resources (L-PBC).

The first year focused on defining S-PBC distribution test cases, identifying operating parameters and constraints for devices for L-PBC, and testing at least one L-PBC algorithm for voltage magnitude compensation on a single distribution feeder. The goal for the end of the first year was to validate PBC as a feasible method for controlling multiple resources on a single distribution feeder. In fact, the project developed and tested *three* local controller designs, each of which was applied to different test cases and scenarios.

In the second year, the goal was to continually refine control algorithms, increase the complexity and scale of the test scenarios, and to prepare all components for the HIL testing and simulation of the PBC paradigm in physical implementation. This included increasing the number of virtually connected nodes in the simulation, tying together the local and supervisory controllers, and implementing the local controller on actual hardware. In the third year, the team's effort was primarily directed at running simulations and comparing performance *in silico* and in the HIL setting.

The objectives of algorithm refinement were to demonstrate convergence under a broad range of scenarios, and to establish robustness in the face of bad data and contingencies. In HIL testing, we were able to establish that no adverse effects resulted from interrupting signal to the local controller, and that the system could recover from ill-chosen phasor targets.

The main hypothesis underlying the PBC control logic is that it is possible to prioritize local physical constraints over system optimzation based on economic criteria, and that such an approach will enhance grid reliability. By acting to restore the phasor target, the local controller will always drive the system toward an operating state known to be safe and stable, without

knowing the nature of the disturbance.[1] Fundamentally, this is because physical operating constraints of the a.c. grid are directly expressible in terms of phasor differences between nodes. The system stabilization achieved by L-PBC then buys time for S-PBC to re-optimize, or update phasor targets with newly available information. This temporal disaggregation – maintaining a technically feasible operating state first, and then adjusting targets to fine-tune the system – inherently prioritizes grid stability and reliability, rather than optimizing independently and then checking for possible violations after the fact.

In the context of deploying distributed solar PV generation, the goal for PBC is to supervise variable resources in real-time to ensure that they will not cause any violation of operating constraints (e.g., voltage or power flow) – precisely the kind of violation which, if found in simulation during an interconnection study, would prohibit adding more PV capacity and imply that the PV feeder hosting capacity has been exhausted. When it can be guaranteed, at the expense of some short-term curtailment, that PV resources will not act to violate constraints because they are tracking phasor targets, the hosting capacity becomes moot and a feeder can accommodate arbitrary penetration levels of solar PV. The project successfully established that PBC allowed the seamless accommodation of >100% penetration.

An overarching design criterion – analogous to medicine – is that phasor-based control should "First, do no harm." This criterion was met throughout. For example, if the local controller could not achieve the voltage phasor target at the performance node, the actual voltage would fall somewhere between the original and the target value. If data flow was interrupted, the controller would default to leaving the actuator with its most recent power injections. In no case did the controller drive the state of the system farther away from the target, or from within the operating envelope to outside the operating limits. Even in simulations where a PBC controller failed to meet performance criteria, the worst-case outcome was a circuit condition equivalent to having no control at all. This was important to ascertain in advance of a field demonstration.

In the first project year, the team specified performance criteria that would constitute success for the control scheme (details presented in Milestone 1.2 Report). The first criterion was nodal voltage control on a distribution feeder to within 0.01 per-unit and 0.5 degrees of angle, in the presence of PV penetrations of 100% of peak coincident load. This was achieved in many simulations under a variety of test scenarios, both *in silico* and with HIL – with the exception of several specific problematic scenarios that prompted further research to define edge cases where controllers failed to converge (e.g., certain uncontrollable nodes or certain actuator locations on a given feeder). Criteria also included a local controller settling time of 10 time steps, or 5 seconds for a 0.5 second time step. Controller settling in less than 10 time steps was routine performance for all three local controller types, again with the above exceptions. The team was not able to implement a 0.5-second time step in the HIL testing environment due to non-negotiable constraints of the inverter control interface (unrelated to PBC).

1 A detailed discussion of the design philosophy with support for this claim is presented in A. von Meier, E. Ratnam, K. Brady, K. Moffat and J. Swartz, "Phasor-Based Control for Scalable Integration of Variable Energy Resources." *Energies* 2020, 13(1), 190.

Controller Implementation and Data Infrastructure

The objective of this task group was to prepare the local controller on a physical software control platform for HIL testing, interface with specific devices including three single-phase inverters. The goal was to show that PBC works in physical practice, with a real controller and real devices, and that it has a viable path to commercialization.

This task group originally envisioned implementing the local PBC controller on an existing state-of-the-art software control platform, designed and built by original project partner 1EnergySystems. Following the acquisition of 1EnergySystems by Doosan Group to become Doosan GridTech, and a subsequent change in management and strategic priorities, Doosan GridTech withdrew from the project at the end of the first year.

The original design objective for the control platform was to provide consistent management and optimization across multiple distributed resources, running operating modes and bulk power applications that help utilities maintain power quality on circuits impacted by solar and other renewables, while taking full advantage of other distributed energy resources such as demand response. The platform was to be built on the open MESA standard to ensure interoperability among energy storage components and cost effective integration of distributed resources with existing grid control systems such as SCADA.

Rather than identifying an alternate controller hardware vendor, the project team chose to build our own. Although we did not package it into a market-ready product, we created a highly functional platform called Distributed Extensible Grid Control (DEGC), building on experience with the eXtensible Building Operating System (XBOS) and the Wide Area Verified Exchange (WAVE) decentralized authorization protocol developed at UC Berkeley. We successfully built and used this new platform in HIL testing of PBC.

This task group also built on the µPMU data infrastructure developed at UC Berkeley under a prior ARPA-E project, including µPMU hardware and the Berkeley Tree Database (BTrDB). By incorporating live data from physical sensors along with controller hardware and devices, the project aimed to demonstrate the practical feasibility of PBC without the need for a full field implementation, which would have been unrealistic under the given budget constraints. High-resolution µPMU data at 120 Hz informed both the design and validation of control algorithms. Streaming µPMU data from the project were hosted and archived for offline analysis on PingThings' PredictiveGrid™ platform, which uses BTrDB technology for managing time-series data streams.

Given the project team's ability to leverage these enabling technologies, the departure of Doosan GridTech did not adversely affect the major project outcome, to demonstrate PBC. It does mean that we do not have a well-established private sector partner to whom the research team can hand off commercialization of the technology. On the other hand, this change resulted in new technology development beyond the original project scope. Based on the demonstrated performance of DEGC and the absence of similarly flexible solutions on the market – particularly based on open-source code – the project team believes that the further development and commercialization of DEGC presents a significant positive opportunity, and a consequential step toward commercializing PBC on a broader scale than might have been achieved in the context of a highly specialized vendor product. In simple terms, we are going longer and bigger than originally envisioned.

Hardware in the Loop (HIL) Testing and Simulation

The objective of this task group was to simulate and then test PBC in a physical facility with hardware-in-the-loop devices, and demonstrate the PBC control paradigm works as predicted in theory. Successful completion of this task will prove the potential of PBC as a realistic and effective new strategy for operating electric grids with arbitrarily high solar penetration levels. The hierarchical PBC structure and its specific algorithms was validated in HIL simulation at the LBNL FLEXLAB.

In this setting, the supervisory controller received state information from the simulated network along with external inputs such as economic and meteorological information and set phasor targets for the controlled nodes on the grid to optimize network operation. Local controllers then modulated the real and reactive power of connected devices such as PV and storage inverters along with controllable loads to maintain the target phasors at the controlled nodes. By constrast with the pure *in silico* simulations, the HIL tests captured latencies that can cause instabilities in the closed loop, including those induced by the communication layer, and limitations of the actual controllability of real devices.

Beyond demonstrating the feasibility and impact of PBC, these tasks aimed to advance simulation technology relevant to the FOA. The team used OPAL-RT's phasor-domain transient stability simulation tool, ePHASORsim, and worked to expand the size of the network that can be simulated.

The first year of the HIL testing and simulation effort focused on drafting the testing protocol and performance criteria for HIL simulations. The second year task was to create plans to install, test and commission HIL devices at FLEXLAB, including the communication and control infrastructure. The goal was to perform all simulations at the FLEXLAB in the third year. The project team actually succeeded in performing some of the simulations ahead of schedule, toward the end of the second year. Overall, simulations showed that most of the original performance targets could be met, speficially those associated with control algorithm convergence, while other criteria such as the length of control time step were limited by the physical equipment.

Phasor tracking was demonstrated in HIL testing to meet the criterion of voltage magnitude convergence to within 0.01 per-unit. The settling time in these tests was 15 seconds instead of the posited goal of 5 seconds, because we had to deliberately slow down the cadence of the controller to agree with the inverter response time, which was in turn limited by the manufacturer's firmware. This does not reflect an intrinsic weakness of the PBC controllers, but a compatibility issue between a high-resolution control paradigm like PBC and today's standard practice by device manufacturers.

Note, however, the linear quadratic (LQR) local controller generally settles with only a single time step iteration, which could be faster that our particular HIL equipment permitted. To prove this, CIL tests were conducted on a 5-second cadence with 0.01 p.u. tracking, meeting the performance criterion.

PBC was also shown robust with respect to missing and bad data inputs. Specfically, we demonstrated smooth recovery when the local controller finds that it is infeasible to attain the phasor target – which could result from a poorly calculated target by the supervisory controller, from a phasor measurement discrepancy, or from a discrepancy between expected and actual

available actuation resource, causing the controller to saturate. This problem was solved by the local controller sending an “I Can't Do It” (ICDI) signal to the supervisory controller, which responds with an adjusted phasor target. Recurring or inconsolable ICDI signals allow for flagging potentially corrupt data. In the event of missing data such as a disconnected µPMU source, the local controller simply fails safe to allow default operation of the distributed resource.

One important lesson learned was the importance of access through the inverter API in order to implement explicit real and reactive power commands used in PBC. The original hypothesis was that inverters could readily respond directly to external signals to increase or decrease net real or reactive power injections. However, even in an R&D setting such as the FLEXLAB, inverter manufacturers tend to be protective of their internal control algorithms, and tend to limit the user’s ability to make detailed adjustments. In a typical commercial installation, for example, a user might just choose among several operating mode settings, which then govern power dispatch in real-time according to some set of criteria. While the inverters used in this project already allowed for explicit real power control – albeit at a slower time step than desired – modifications had to be made to allow for explicit, independent reactive power control, as opposed to changing a power factor set point. This is not technically difficult, but a question of policy. Consequently, for PBC to recruit broad participation from DER, some cooperation on the part of inverter manufacturers will be required.

General Findings

In sum, we have established that PBC works effectively to track target voltage phasors to within a narrow tolerance of 0.01 per-unit by recruiting real and reactive power contributions from DER, under a broad range of scenarios and with high levels of PV penetration. The only discrepancy between the system performance in HIL testing and the original performance objectives is that the settling time for inverter control is longer, owing to a longer time step mandated by the inverter hardware/firmware.

We also completed a comprehensive general survey of PBC stability and feasibility, beyond the simulations on realistic test circuits, to better understand the limitations of the method. In sum, we found the factors that were originally considered as likely candidates to challenge PBC – controlling many nodes on large circuits, with high penetration levels of PV – not to be especially problematic. Difficulties encountered with large circuits mainly had to do with proper impedance modeling, not with the number of nodes. Though we now have a very good idea how to contrive hypothetical scenarios that will defeat PBC, none of these are likely to be encountered in the field.

We also have demonstrated an effective data infrastructure for PBC, where µPMUs, supervisory and local controllers as well as loads and inverters communicate and cooperate. We have generalized this as an open-source Platform for Distributed, Extensible Grid Control (DEGC) that can serve for other control approaches as well. While this platform proved to be very reliable in testing, the PBC paradigm has also been shown to be robust with respect to missing or incorrect data.

Finally, the team completed a successful value analysis for PBC at the transmission level. The selection of use cases for simulation and testing within the scope of this project deliberately focused on applications of PBC that could be primarily motivated by distribution-level

considerations, since it was necessary to establish a technical proof-of-concept at the distribution scale first: without this, transmission level applications would not be technically feasible. However, with input from PG&E and Dominion Energy, the team determined that the most *economically* compelling early use cases for PBC would likely be motivated by grid security considerations, rather than cost-benefit analyses at the scale of individual distribution feeders. Although it is strategically important that PBC supports very high penetration levels of variable solar PV generation, we take this as a prerequisite rather than an explicit selling point, since it is difficult to assign an economic value to the utility from increasing feeder hosting capacity. Similarly, the ability to improve power quality on distribution circuits suffering from high volatility in either generation or load is hard to monetize directly in the distribution context. By contrast, extrapolation of the impact of PBC-controlled distribution circuits on the transmission level captures tangible economic benefits. Using value of lost load (VOLL) and timing on equity analyses for scenarios with a single planned system upgrade, we found that mitigation of security constraints by PBC in the scenarios studied can yield a present value on the order of $1 million, accomplished by a single, well-placed PBC performance node with a hardware cost on the order of $10k.

Based on these results, we believe the project has successfully established PBC as a feasible technology, ready to advance to field demonstrations and next steps for commercialization.

Project Activities

Italicized subtask descriptions below are taken directly from the SOPO. They are followed where applicable by discussion of original hypotheses, approaches used, problems encountered and departures from planned methodology, and assessment of their impacts on project results.

Case Design

Subtask 1.1 *Build Test Models (Month 1-Month 9): Select sample circuits (including well-understood IEEE test feeders and realistic, representative, anonymized, non-confidential circuits) and build test circuit models for PBC, which may include changing format and editing details, so that simulations can be run on OPAL-RT. Include specifications for DER and loads, where the proof-of-concept feeder-level simulation of PBC will not require very particular load models. The test models will be placed and maintained in a private grid model repository set up specifically for the project leveraging technology built by GridBright for the ARPA-E Grid Data project. Where prior ownership rights allow it, the models will be shared with the IEEE for inclusion in the IEEE test library with the intent that they will be published at the end of the project (no longer than 12 months past the completion of the project) after model updates have stabilized.*

Circuit Models

The original hypothesis was that among such a diversity of circuit models, some challenges for PBC would be discovered that would help to identify the range and scope of applicability of the PBC paradigm. This proved to be true. Some circuits and scenarios posed particular challenges to local controllers, sometimes in ways that were not obvious. For example, it stands to reason that a single-input, single-output (SISO), proportional-integral (PI) local PBC controller (which decouples real power and voltage phase angle from reactive power and voltage magnitude) would be less effective on circuit with a low X/R ratio (where decoupling is physically less prominent). However, it was also found that certain specific placements of multiple control and actuation nodes on some of the larger circuits did not yield satisfactory convergence of the PI controller. These observations motivated the creation of a special tool to identify controller feasibility regions, or “heat maps,” which can predict whether local PBC would be challenged by a particular configuration and recommend effective actuator placements – which, in practice, would translate into preferred siting of controllable DER.

The GridBright infrastructure for publishing circuit models created through the ARPA-E Grid Data project proved to be useful. Models were uploaded to the repository and edited throughout the project as they became ready.

The conversion of circuit models proved to be even more time consuming than anticipated, despite the project team’s expectation that this would entail significant effort and some degree of frustration. The number of ways in which small transcription errors or inconsistencies can derail a model conversion process is astonishing. Project team members from GridBright, UC Berkeley and OPAL-RT all collaborated on debugging, often with multiple iterations, as errors became apparent in the process of attempting to run simulations. The milestone of 30 validated,

published circuit models was met despite setbacks, and they can be found on BetterGrids.org. The most ambitious of these – namely, the 13 circuit models based on PG&E Taxonomy Feeders – are summarized in Table 1 below.

No.	Circuit model	# of Buses	# of Voltage Regulators	Notes
1	AL0001	1139	0	76 mile 12kv feeder
2	AT0001	691	0	36 mile 12kv feeder
3	BR0015	546	0	8 mile 12kv feeder
4	BU0001	167	0	7 mile 4kv feeder
5	DO0001	1487	0	17 mile 12kv feeder
6	HL0004	2123	3	331 mile 21kv feeder
7	MC0001	1147	6	202 mile 12kv feeder
8	MC0006	1747	5	145 mile 12kv feeder
9	MO0001	2413	0	51 mile 12kv feeder
10	OC0001	1116	5	106 mile 12kv feeder
11	PL0001	344	0	20 mile 12kv feeder
12	R1-12.47-1	1828	0	30 mile 7kv feeder
13	TMP0009	4838	8	533 mile 12kv feeder

Table 1: PG&E Taxonomy Feeders

PV and Load Data

The volatility inherent in high-resolution (such as one-minute) data is important for simulating control behavior; however, many available datasets use 15-minute or longer intervals. Also, the goal was to simulate very high penetrations of distributed solar PV on these circuits, which is of course not part of the standard models. Thus, to be useful for this project, distribution circuit models had to be populated with customized generation and load data.

The team investigated several potential data sources, including Pecan Street (load data) and NREL (PV data). Unfortunately, the licensing rules for Pecan Street the data would have precluded making our completed models publically available. Consequently, the team opted to generate synthetic data based on public, residential and commercial load profiles from Southern California Edison. To augment the common, representative variations of loads and PV generation, we added aggressive step changes (e.g. an instantaneous change of 50% of total feeder load) to challenge the controllers.

Subtask 1.2: Define Objectives (M6-M12): Define constraints, objectives and optimization criteria for four PBC test cases (such as feeder loss minimization, voltage magnitude regulation, and real/reactive power line flows regulation, and any other as deemed necessary). Define performance metrics for PBC, such as response time, tolerance for tracking the phasor target, and robustness of control algorithms under varying conditions with multiple participating DER and any other as deemed necessary.

A workshop with utility experts (from PG&E and Dominion Power) was held to define control objectives with relevant use cases for PBC in the context of distribution circuits. The four control objectives were:

1. Phasor alignment, i.e. matching an externally given phasor to support switching operations;
2. Phase (ABC) balancing;
3. Voltage volatility mitigation; and
4. Maintaining a target voltage profile.

Note that the target voltage profile might include voltage magnitude only, in which case it would be designed to narrow the operating envelope for efficiency purposes, such as Conservation Voltage Reduction (CVR). If including both magnitude and angle, however, the voltage phasor profile amounts to a complete specification of the operating state, corresponding to a specific real and reactive power flow between the performance node and the reference (given a constant, known complex impedance between these two points).

The team defined performance metrics for latency, settling time, steady-state tracking error, and recovery from missing or infeasible phasor targets. The first criterion was nodal voltage control on a distribution feeder to within a tolerance 0.01 per-unit and 0.5 degrees of angle. Criteria also included a local controller settling time of 10 time steps, or 5 seconds for a 0.5 second time step. Over-and undershoot limits of 6% and 13%, respectively, were chosen on the basis of the IEEE 1547 Standard. Robustness of control algorithms was defined in terms of satisfactory performance under loss of communications, controller saturation, and discontinuities such as topology changes on the circuit. In no case should the controller action result in an operating condition outside the permissible envelope, or farther from the intended target than without control.

Subtask 1.3: Scale-up Models (M3-M12): Collect and anonymize 30 realistic candidate distribution circuit models for use in scale-up simulation. Of these 30 models, at least 20 should prove suitable for demonstrating different features and capabilities of PBC. In cases where the model authors are not interested in publishing updated models and the original ownership rights allow the models to be publicly published, the project team will place the model on a permanent public repository, assign new DOIs and cite the models using the DOIs in all IEEE publications.

This subtask was completed. All of the models turned out to be usable in some way, and all could be published on the BetterGrids.org repository website.

Subtask 1.4: Cybersecurity and Interoperability Plans (M1-M3): Produce the initial cybersecurity and interoperability plans for PBC platform.

This subtask was completed, led by GridBright. The first iteration addressed the following questions: How the various components in a working PBC system will communicate with each other, and in what format; how the data archive in BTrDB should relate to the components directly involved in the closed-loop control architecture; which aspects of information from the Traditional System Layer, if any, are initially important to consider; and what concerns, if any,

arise regarding compatibility between and among common data protocols used in systems with which PBC might interact. The team did not identify any substantial security or compatibility concerns.

Subtask 1.5: Market Transformation Plan (M9-M12): Draft a market transformation plan that includes product development, architecture designs, identification of target market, and legal/regulatory considerations including intellectual property, infrastructure requirements and data dissemination. The project team will make control algorithms available via journal publications, and, to the greatest extent possible while respecting utility confidentiality and ownership rights, will strive to publish test feeder models in a permanent public repository.

This subtask was completed, led by Doosan GridTech. The market transformation strategy shifted in subsequent quarters, as the original hypothesis was that PBC would become a functionality of an existing battery controller, and Doosan GridTech would act as the driver for commercialization. When Doosan left the project after Year 1 due to a change of management post-acquisition, the project shifted toward a more fully open-source approach, centering on the Distributed Extensible Grid Control (DEGC) platform developed at UC Berkeley which is detailed further below.

The impact on project outcomes was twofold. On the one hand, not having PBC installed on an already commercialized product will constrain near-term opportunities for physical deployment and thus delay the introduction to the industry, potentially by several years. On the other hand, with DEGC the project created a valuable, unexpected product beyond the original scope. Specifically, we built an enabling technology: a secure platform that supports the participation of independent DER owners and operators in response to physical grid measurements, without compromising privacy and without divulging more than a bare minimum of sensitive grid operating information (e.g., one specific μPMU measurement as a reference value non-local to the controlled resource).

The original hypothesis was that a utility would deploy PBC in conjunction with their own assets, such as a substation-scale battery, as part of a controller package that happens to include PBC based algorithms as a subset of controller options and settings. The new capability enabled by DEGC is the secure recruitment and coordination of DER owned by multiple, independent parties. While it will take longer to realize this in a field deployment, it is more consistent with the original vision for PBC as a common framework for recruiting many diverse resources.

Subtask 1.6: The team will develop and submit to DOE an agreement on intellectual property (IP) framework, IP Agreement Plan, with industry partners that covers pre-existing IP, ownership of IP (including definition for what constitutes contribution to a given work), patenting rights, exclusive and non-exclusive licenses, and settlement of disputes.

This subtask was completed. Note that the early IP framework reflected the original assumptions about Doosan GridTech's role on the commercialization team. The project subsequently moved toward a more open-source approach.

Subtask 5.3: Scale-up Models (M13-M24): Gather and build additional circuit models for OPAL-RT simulation to support diverse and challenging scenarios for PBC. For the purpose of demonstrating scalability, some circuits may be duplicated in the simulation as needed. The goal is to perform as large a simulation as is realistically achievable within the project budget, to demonstrate that the scaling of PBC to a larger number of interconnected feeders, and the OPAL-RT simulation thereof, can in principle be extended using the proven methodology to encompass 1,000,000 nodes as defined in the ENERGISE FOA.

The OPAL-RT team spent considerable effort on the scale-up, beginning with the creation of a synthetic transmission and distribution interconnection model. For this purpose, 50 distribution circuits were connected to the standard IEEE 118-bus transmission network model. The combined model consisted of 185,000 nodes and was run successfully on a real-time machine OP5700 utilizing 16 CPU cores, at time-step of 10 ms.

In the interconnected simulation, the distribution model was treated as a lumped load for the transmission system analysis, whereas the transmission network was seen as a voltage source in distribution analysis. A dynamic Thévenin equivalent of the transmission network was used in the distribution model to replace the substation voltage sources. The Thévenin equivalents were updated in terms of series impedance and a voltage source behind them, to deliver the required power to the feeder and maintain the voltage at the substation bus. Conversely, in the transmission model, the distribution systems were represented as variable impedance loads whose value was adjusted based on the power consumption at the distribution substations.

The power consumptions of the distribution feeders were feedback to the respective lumped loads at IEEE-118 network. Similarly, respective Thévenin voltage and impedances from IEEE 118 network were considered as voltage sources to the distribution feeders. To confirm the successful interconnection of the transmission and distribution systems, two case studies were performed to analyze the propagation of disturbances through the interconnected system. First, a three-phase ground fault on the transmission system was simulated, and its effects observed on the distribution system; this resulted in voltage sags varying with distance along the distribution feeders as expected. Second, a step change in load at the distribution level was observed to have expected impacts at the bulk level, as evidence by generator swing behavior. This effort substantially advanced the state-of-the-art in dynamic system modeling.

Figure 2 represents the overview methodology of the interconnection of the transmission and distribution feeders in the ePHASORsim model.

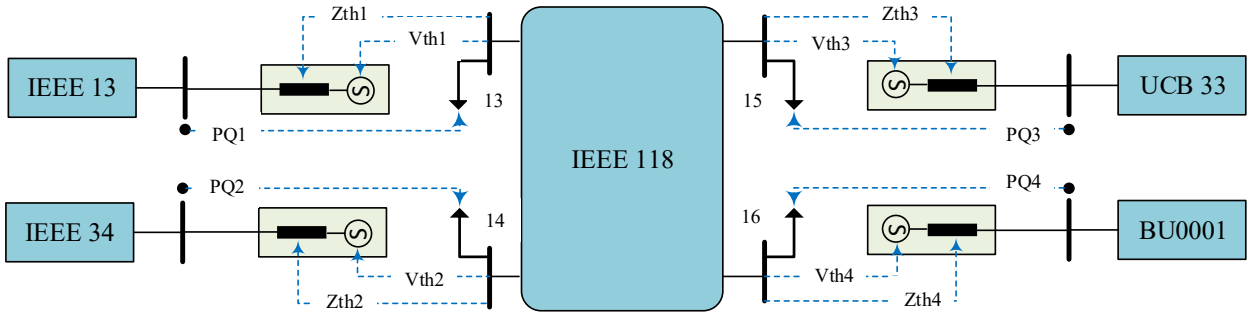

Figure 2: Overview of transmission and distribution feeders in OPAL-RT's interconnected model.

Control Algorithm Development

Subtask 2.1.1: Define S-PBC (M1-M9): Develop algorithms for four 'Supervisory Phasor-Based Control' (S-PBC) distribution feeder test cases based on objectives identified in Subtask 1.2. The outcome of this subtask is, for each test case, a set of target phasors for two or more nodes on the feeder intended to produce the particular optimization outcome.

Supervisory PBC was successfully formulated to generate phasor targets for each type of objective. For example, the objective of A-B-C voltage magnitude balancing at all nodes was incorporated in a day-ahead constrained optimization problem, solved at the supervisory layer. The constraints include a linearized model of the modified IEEE 13-node test feeder. The output of the constrained optimization problem are day-ahead voltage phasor targets, which are communicated to the local control layer. The local controller then tracks the day-ahead voltage phasor targets, while rejecting disturbances and responding to solar PV generation and load variability.

The performance of S-PBC is not closely tied to the particular nature of the chosen objective, or the specific information fed to the optimization (e.g., generation and load forecasts or cost functions). In that sense, S-PBC is agnostic to the solar PV penetration level of the system. Given enough computing power, S-PBC is also agnostic to the time step on which it is optimizing (e.g., day-ahead, hour-ahead or minute-ahead.) To be practical, though, S-PBC should be able to run on modest computational resources like a standard laptop, on a time scale of minutes – that is, a time scale where it is reasonable to expect a re-optimization based on contingency events or changes in the forecast.

The basic challenge for S-PBC therefore lies in applying a suitable linearization of the power flow equations that does not introduce excessive errors, but is computationally fast. Intrinsically, the problem of expressing power flows in terms of voltage phasor differences is poorly conditioned, in that small errors in phasors correspond to large discrepancies in power flows. This first round of S-PBC algorithms worked, but was too slow to promise effective scaling to

larger networks. Over the course of the project, several different approaches to the linearization were employed, and some newly developed.

Linearization

We knew from the outset that the key challenge for S-PBC would lie in obtaining an effective compromise between computation time for an optimization, and fidelity to the nonlinear the power flow equations that relate voltage phasors to nodal real and reactive power injections. However, the effectiveness of any compromise had to be tested empirically. To this end, the project team developed or adapted several linearization approaches, including a Loss-Approximated Power Flow (LAPF) and an iterative linearization technique.

Much of the recent research into OPF solutions has focused on convex relaxation techniques, where 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 semidefinite 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.[2] 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 looked 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 established relationships between the active and reactive power flowing through conducting lines and the squared voltage magnitudes at each of a radial network's nodes.[3] 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. The Simplified DistFlow equations were later generalized from their original, single-phase form to unbalanced, multiphase networks. Building on this three-phase linear model, the work of Sankur developed the LUPFM, introducing an additional relationship between voltage phasor angles and power flows into a model that was previously only capable of specifying voltage magnitudes.[4] This allowed for the treatment of a phasor in its entirety within the structure of the model.

The main benefit of defining our OPF as a QP using the LUPFM is that it is a very dependable means of generating a feasible set of voltage phasors. Its disadvantage is that the linearization of the power flow equations necessarily introduces an approximation error that prevents the end solution from being feasible with respect to the original power flow equations. To address that issue, we developed an iterative solution method that involves successive exchanges between the

2 A description of some conditions under which convex solvers return exact solutions is given in 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.

3 M. E. Baran and F. F. Wu, "Network reconfiguration in distribution systems for loss reduction and load balancing," in *IEEE Transactions on Power Delivery*, vol. 4, no. 2, pp. 1401-1407, April 1989.

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

LUPFM-based OPF and a nonlinear, Newton-Raphson based solver. At each step of that iterative method, the nonlinear solver is used to refine the LUPFM, allowing the end solution of the joint method to approach an exact solution of the power flow equations.

Subtask 2.1.2: Test S-PBC (M8-M12): For each case defined in Subtask 2.1.1, check the correspondence of phasor targets and power flows with DER on a test feeder. The outcome of this subtask is evidence that the target phasors as specified for each test case produce the desired optimization outcome, in the steady state.

The team established that phasor targets produce power flow results to within the required tolerances. Milestone 2.1.2 in Year 1 demonstrated agreement of real and reactive power values as dispatched by S-PBC to within 5% of the desired steady-state values computed by an Optimal Power Flow (OPF), which was the originally stipulated criterion.

As the first approach to achieving this metric, the team used a full nonlinear power flow solver, OpenDSS, and integrated it with the linearized optimization. With this OpenDSS integration, the structure of the S-PBC code was as follows:

1. Import an ePHASORsim model and load profile;
2. Convert to a linearized unbalanced power flow model and solve an optimization with respect to a feeder-level objective;
3. Export the resulting real and reactive power dispatches at each controllable DER to OpenDSS and solve the feeder's nonlinear power flow equations;
4. Report the resulting voltage phasors to the L-PBCs as targets.

While this method worked, it also became clear that it would be too computation intensive to scale well. Therefore, subsequent efforts focused on improving the linearization and approximation procedures.

Subtask 2.2.1: Define L-PBC (M1-M12): Identify operating parameters and constraints of devices for Local PBC (such as, achievable real and reactive power, response times, switching constraints for loads and DER) relevant to objectives and test cases defined in Subtask 1.2, and develop an algorithm for recruiting two or more resources to track a given phasor target.

The project team pursued three distinct L-PBC controller designs with different degrees of complexity to achieve the various feeder-level control objectives specified in the test cases. They are (1) a simple Proportional-Integral (PI) controller, (2) a Linear Quadratic Regulator (LQR), more specifically described as a Linear Quadratic Tracking Controller (LQTC), and (3) a Retrospective Cost Adaptive Controller (RCAC). (1) and (2) were both developed by UC Berkeley, and (3) by the University of Michigan. Each controller is designed to track a voltage phasor target while rejecting disturbances, by way of modulating real and reactive power (P and Q) commands for one or more distributed resources.

The PI controller is effectively a combination of two single-input, single-output (SISO) controllers, one for voltage magnitude and one for angle. Both the LQR and RCAC are multiple-input, multiple-output (MIMO) controllers. This distinction is illustrated in Fig. 3.

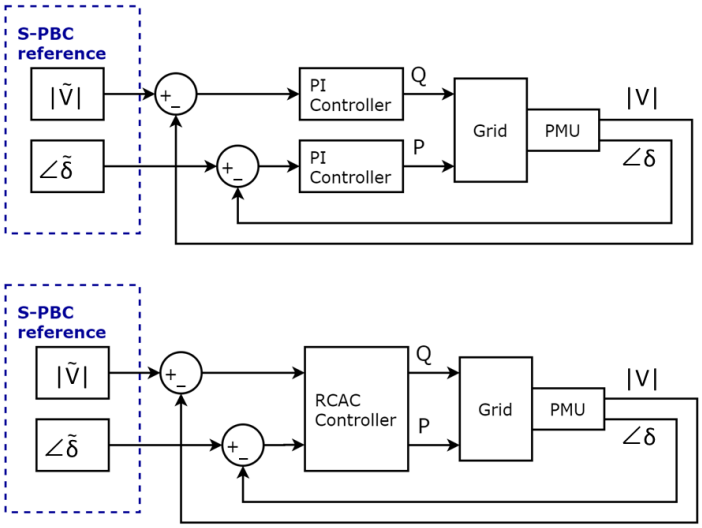

Figure 3: Block diagram for the SISO (top) and MIMO (bottom) control architectures.

A key design challenge for each of the local controllers is addressing the coupling between different L-PBC nodes and ensuring interconnected stability. Here, "coupling" refers to the fact that the real and reactive power (P and Q) actions at one L-PBC node will unintentionally affect the voltage phasors (states) of other L-PBC nodes on the same circuit. This coupling depends on the network characteristics and operating state of the system and is unknown to each local controller. It is an established result of systems control theory that the interconnection of stable sub-systems does not necessarily result in a stable interconnected system. Thus, while an L-PBC algorithm may track a phasor target perfectly at a single feeder node, it may be disturbed by the presence and activity of other controllers on the same feeder. These interactions are not obvious to predict and have to be studied empirically in simulation scenarios.

"Coupling" can also mean the interdependence between P-delta and Q-V variable pairs. The respecitive sensitivites depend on the circuit impedance, and specifically the ratio of inductive reactance X to resistance R. When X is much greater than R (as is typically the case for transmission systems, but not always for distribution lines), the variable pairs are effectively decoupled. The hypothesis was that as a MIMO controller, the LQTC would be better able to accommodate different sensitivities. However, the LQTC requires an effective impedance estimate, which would typically be obtained from a circuit model. The advantage of the PI controller is that controller gains can be determined empirically, without knowledge of a circuit model.

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. Like the PI controller, RCAC does not rely on any externally provided information about impedances within the network. It is designed to re-optimize the control law over a trailing data window (thus "retrospective") with a recursive least-squares algorithm. A re-optimized controller is implemented at the next time step. This enables the controller to both

improve its performance by learning, and readily adapt to transitions of phasor targets from S-PBC.[5]

The original hypothesis was that different types of controllers would exhibit different advantages or disadvantages in specific situations. This proved correct, although each of the three controllers worked in most scenarios. While RCAC was expected to be the most sophisticated and elegant controller design, the project hedged against the risk that RCAC could not be implemented in the HIL setup by developing a simpler set of algorithms in parallel, so as to guarantee that we could implement local PBC with a very basic and well-understood control approach. This was an effective strategy, as we opted to perform HIL testing with only the PI and LQTC controllers.

Subtask 2.2.2: Test L-PBC (M7-M12): For each of four test cases, test the ability of local PBC with DER to attain the phasor target specified for two or more nodes in a quasi-steady state, with moderate transitions from one target to another. The outcome of this subtask is evidence that the local control algorithm effects tracking of a steady-state phasor target, and without creating instabilities or adverse interactions among resources.

A large variety of test scenarios were run in simulation with all three controllers. These simulations mixed and matched the four control objectives with the different test feeders, different amounts of PV penetration and different placement of actuator and performance nodes on the circuits, as well as the three controllers. Out of the ~10^4 combinatorial possibilities, our team collectively simulated several hundred. Rather than trying random combinations, we initially focused on the IEEE 13-node feeder so as to develop an understanding of what circumstances challenge the L-PBC controller, and to further push in those directions. It should be noted that despite its small size, the 13-node feeder with its single- and two-phase branches presents a very rigorous phase imbalance challenge for the controller, which (due to coupling) becomes relevant for all the control objectives, not just phase balancing. Although L-PBC occasionally could not comply with reaching its target, we established that, fortunately, it is very difficult to create adverse interactions on a circuit with PBC. Across hundreds of simulations, the worst-case outcome was the status quo of power flows without control.

A simple example of local PBC on the IEEE 13-node feeder is illustrated in Fig. 4. Here, actuators consisting of controllable DER at three different nodes share the task of maintaining a constant net power export of the distribution feeder across the substation transformer, represented by the voltage phasor difference between node 650 (reference) and 652 (performance node). Uncontrollable loads and solar PV generation are simulated on other nodes, as indicated in the diagram. Two large step changes in these net loads are quickly mitigated by the actuators, causing the voltage to gently settle back on the prescribed target values (labeled as V_{ref}) on each phase. While Fig. 4 shows only voltage magnitude, the voltage angle settles in an analogous manner.

5 Background on the RCAC controller can be found in Y. Rahman, A. Xie and D. S. Bernstein, "Retrospective Cost Adaptive Control: Pole Placement, Frequency Response, and Connections with LQG Control," in *IEEE Control Systems*, vol. 37, no. 5, pp. 28-69, Oct. 2017.

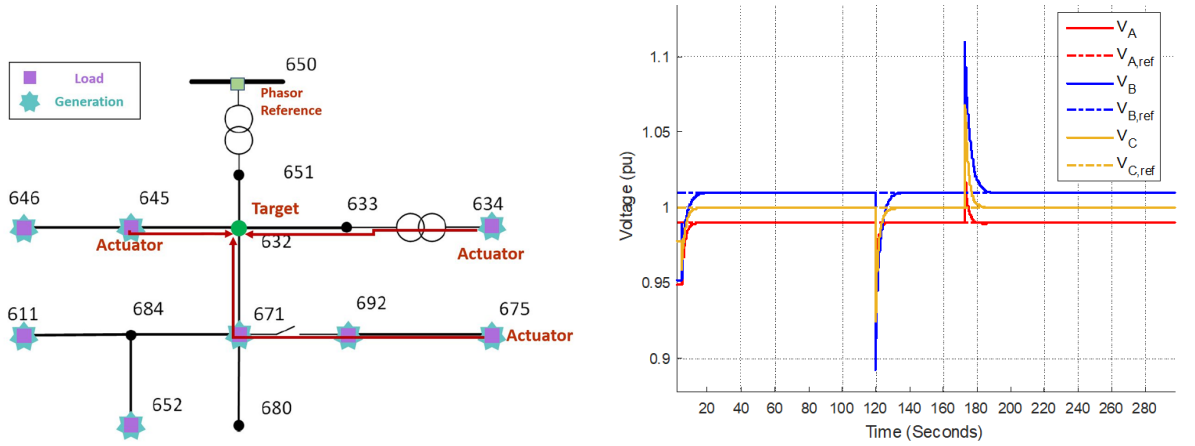

Figure 4: Disturbance rejection by Local PBC.

Subtask 6.1.3: Refine and Test S-PBC (M13-M24): Develop and improve S-PBC algorithms with more challenging scenarios such as larger feeders and phasor targets specified for more nodes. Continue to test how optimization outcomes are achieved through specification of phasor targets along a feeder.

Throughout the course of the project, numerous improvements were made to S-PBC algorithms. Overall, we found that the PBC paradigm proved workable. No problems or constraints were encountered that would be show-stopping for either the simulations within the scope of this project, or for future scale-up. Managing computational effort was the major concern. Since the PBC paradigm specifically allows an arbitrary time step for supervisory control, depending on the urgency and level of sophistication of the economic optimization that is desired, the computing power to throw at the problem is somewhat negotiable. Our S-PBC algorithms were successfully tested on distribution circuits with hundreds of nodes, as well as the 118-bus transmission network, and completed within minutes on an ordinary laptop.

Optimization

With respect to S-PBC optimization objectives, the team decided that simplicity would be a virtue. The important aspect for this project was to translate the constraints and chosen objectives for the circuit into the correct phasor targets, rather than pursuing subtleties of the optimization itself (for example, complicated generator cost functions).

To provide flexibility, we refined a cost function for the S-PBC to implement as either a general stand-alone objective, or in conjunction with primary PBC objectives (target matching, phase balancing, and voltage volatility), with adjustable weights for the primary objective relative to the cost function. For our purposes, it was sufficient to use a simple linear relationship between generation and total cost that allows both positive and negative costs. Our S-PBC framework allows for the future substitution of more sophisticated optimization algorithms.

Because of the central role of solar PV generation for this project, we explicitly accounted for solar capacity. Our solar envelope calculator uses solar geometry and local coordinates to calculate direct beam radiation, to set a real-time (or simulated) maximum PV resource under clear-sky conditions. Combining this with the fraction of solar capacity in a given actuator gives

a final solar weight, which is applied to the actuator capacity as an optimization constraint. This provides the functionality of reducing error between expected capacity, as seen by the S-PBC, and actual capacity, as seen by the L-PBC. By reducing this error, we could reduce the likelihood of reaching a state of saturation for the L-PBC. This framework also allows for future improvement with more sophisticated solar and battery state-of-charge forecasting algorithms.

Impedance Estimation

An important consideration throughout the project was that reliable network models, especially for distribution circuits, may be unavailable in practice. Consequently, both S-PBC and L-PBC must be designed in a way that is robust with respect to missing or inaccurate network impedance information. To this end, online impedance estimation methods were developed.

For L-PBC, the concern is limited to the effective impedance between a given actuator node and the reference node, usually the substation, which determines the effect of a change in power injection on the target voltage phasor difference. Here we may assume that the change in power is entirely absorbed by the substation, which acts as an infinite bus voltage source. L-PBC can then estimate the effective impedance between its local node and the substation by learning the linear relationship between changes in the local phasor voltage measurements and changes in local current injections from the actuators.

By contrast, for purposes of optimization, S-PBC requires a network model that includes the impedances or admittances of every connection. For transmission networks, a good estimate of the bus admittance matrix $\mathbf{Y}_{bus}$ is usually available, but for distribution systems there is often no such thing.

The UC Berkeley team was able to prove that, due to the rank constraints of voltage and current data sets for electrical networks (with at least one node that does not have a power injection or load), the full $\mathbf{Y}$ matrix cannot be estimated from phasor measurements using Ordinary Least Squares or other regression techniques. However, a Kron reduced matrix $\mathbf{Y}_K$ of smaller dimension, which preserves the effective admittances between the leaf nodes of $\mathbf{Y}$, can always be recovered from phasor voltage and current measurements, if all the leaf nodes in the network or segment are measured. Furthermore, if the network (segment) is radial, then $\mathbf{Y}$ is fully recoverable from $\mathbf{Y}_K$. For non-radial networks, $\mathbf{Y}_K$ can be used in the S-PBC OPF calculations because it maintains the input-output properties of the network between the leaf nodes.[6]

With this, we established the general feasibility of PBC even in situations where network models are unavailable.

Application to the Transmission Level

Although the project focused on developing PBC for distribution circuits, the extension to the transmission level was of considerable interest to the research team. Here, the networked topology poses some unique challenges.

6 Details are presented in K. Moffat, M. Bariya and A. von Meier, "Unsupervised Impedance and Topology Estimation of Distribution Networks—Limitations and Tools." *IEEE Transactions on Smart Grid* 2020, 11(1).

For a radial distribution circuit, the substation represents a stiff voltage source, and a logical reference node for defining the phasor. We can then use the effective impedance between an L-PBC node and the substation as a model for how a change in local power injection will change the relative phasor. On a meshed network, it is not obvious which node should be used as the phasor reference node. There is likely not a single node acting as a voltage source; rather, there may be many or zero, with some nodes implementing droop control. Consequently, a change in power injection at any given actuation node could also be met with unintended changes in the power injections elsewhere.

In order to quantify the actual effect of a substation power injection on the phasor relative to a transmission level reference, we started with an effective impedance model, and then added a feedback component above the distribution-S-PBCs that effectively works like an L-PBC to correct the inevitable error. Our transmission network-facing, MIMO, error-integrating feedback controller determines the substation power necessary to maintain the assigned transmission phasor target using feedback control. This substation power is passed into the distribution-level S-PBC, which readily includes the substation power injection in the objective function for the local feeder.

This PBC design paradigm is modular, with multiple fractal layers are possible. Fig. 5 illustrates the case of two layers. Here, a network-facing feedback controller at the substation informs targets relative to transmission system, acting like L-PBC facing upward and as S-PBC downward. The parent supervisory controller is agnostic to whether child node is S-PBC or L-PBC. The successful definition of this modular framework was a crucial proof-of-concept to convince us that PBC will in fact be arbitrarily scalable.

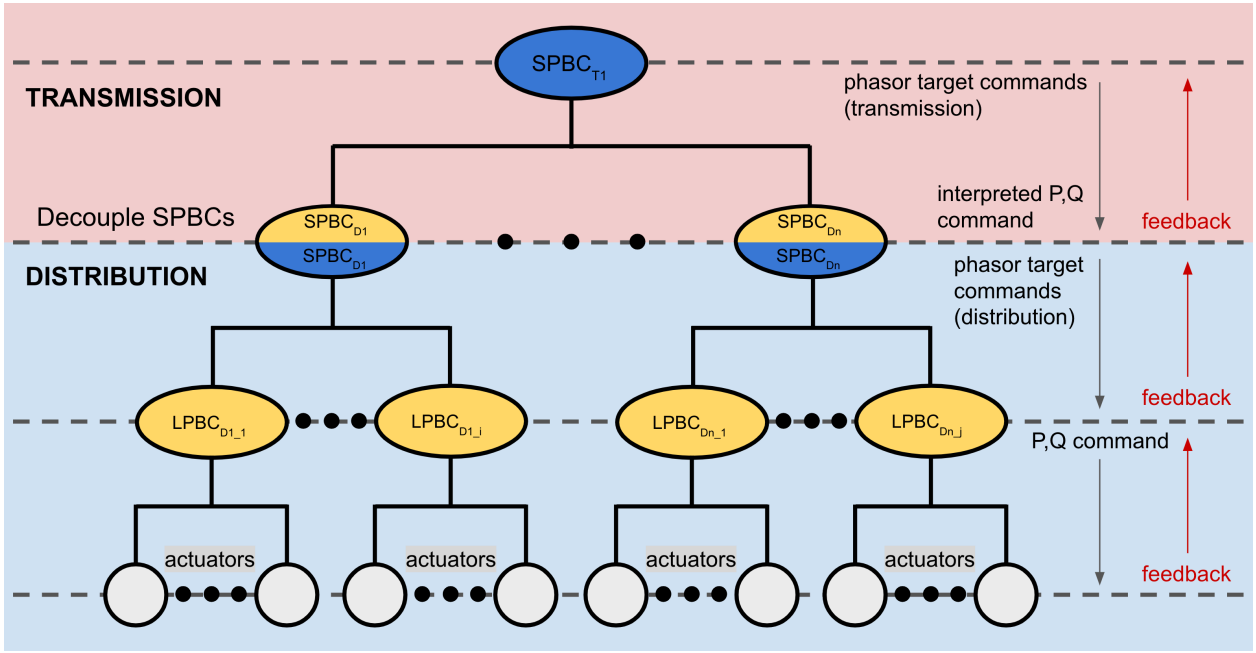

Figure 5: Layering of PBC.

The Berkeley team successfully tested this modular S-PBC approach on a transmission network with 50 interconnected distribution circuits. Decoupling the optimization and power flow processes decreases the computational burden of any individual supervisory controller and creates a problem formulation that is well suited for parallel computation. The transmission level S-PBC assigned target real and reactive power injections at each actuator node, represented by each distribution feeder. These P,Q target injections are interpreted by the distribution level S-PBC as one optimization objective, paired with a local objective that is concerned with conditions on the distribution feeder and relative contributions of DERs. Each objective is multiplied by a respective weight to allow for tuning.

Subtask 6.1.4: Handoff S-PBC (M19-M21): Prepare S-PBC algorithms for implementation on a virtual optimization platform, to produce a supervisory control signal in response to input data (e.g. system disturbance, cost functions and constraints).

The original hypothesis was that PBC would be implemented as one of many functionalities on an existing commercial DER controller. Instead, the project developed its own DEGC platform, detailed below, which supports all information flows between grid measurements, S-PBC and L-PBC. The S-PBC optimization was run by the Berkeley team on a standard MacBook laptop computer, with no handoff required.

Subtask 6.2.3: Refine and Test L-PBC (M13-M24): Develop and improve L-PBC algorithms (e.g. adaptive algorithms) with more challenging scenarios (e.g. increasing numbers of feeder nodes, L-PBC controllers on a circuit, and DERs under a single controller). Introduce transient disturbances (e.g. device outages, voltage sags, faults) into test cases. Test the ability of the local controller with DER to track and restore the phasor target when transient disturbances are introduced, without adverse interactions among resources.

Each of the three local controller approaches were tested in a variety of scenarios and continually refined throughout the project.

PI Controller Tuning

The key element in developing a PI controller is to tune its gains (i.e., how aggressively the controller responds to the feedback signal). The Berkeley team developed two alternate methods that were both successfully tested. The first is the traditional Ziegler-Nichols method. While conceptually straightforward, this method has several practical disadvantages: it must be performed online with actual grid measurements; it involves manual tuning that pushes the real system to a potentially dangerous point of marginal stability; and it needs to be re-tuned regularly if parameters change.

An alternative tuning method allows offline calculation of PI controller gains. The inputs are transient step response requirements, and a network model including time-varying or dynamic components. The method uses a genetic algorithm to cycle through many controller designs, to find one that meets the transient response requirements. This involves pole placement techniques for grid dynamics that can be approximated as first or second-order processes. It incorporates

step response data, simulated in RT-Lab, for each individual actuator and phase, thereby accounting for unbalanced systems.

The genetic algorithm tuning method is faster to execute and generally preferable to online tuning for the reasons above. However, to be conducted fully offline, it requires an accurate network model. The Ziegler-Nichols method provides a hedge against the risk that no such model is available. With this redundant approach, we gained confidence that the PI controller – the most straightforward of L-PBC designs – can be effectively used in most control scenarios. There remain situations, however, for which no controller gains could be found that would yield proper convergence.

Fig. 6 illustrates the improvement in settling time from manual controller tuning (left) by the genetic algorithm (right), for a simple test case on the IEEE 13-node feeder. The control objective here is to drive the voltage magnitude to 1.00 per-unit at the performance node using a DER actuator at another node, and to maintain that target in the face of step disturbances.

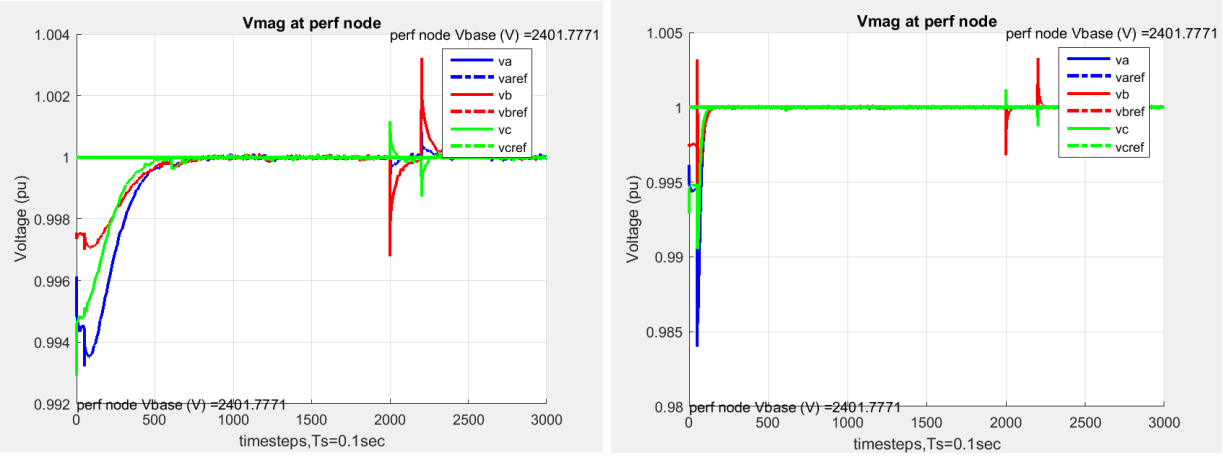

Figure 6: PI controller tuning with genetic algorithm.

PI Controller Placement Tool

The team systematically explored the limits of the PI controller's capability. For this purpose, a tool for assessing the stability of PBC configurations was developed to determine whether for a proposed configuration of actuators and performance nodes, there exists a set of controller gain parameters that will ensure the local controllers effectively track the voltage phasor target. This tool runs on large feeders to identify good actuator-performance node configurations, so that the S-PBC can then set these configurations to determine optimal phasor targets. One challenge was to reduce the computation intensity when moving from single-phase to 3-phase implementation.

The stability assessment approach requires solving a nonlinear semidefinite program. It was determined that the off-the-shelf solver PENLAB was inaccurate at solving this problem on small networks, and because the problem is NP-hard, it was clear the problem would not scale well. The team decided to use a different method, where we sample over a controller gain parameter space and evaluate whether the closed loop eigenvalues are stable. Thus a configuration of actuator-performance node pairs is said to be "feasible" if we are able to find a

set of SISO integrator controller gain parameters that results in performance node voltages reaching their phasor targets. However, we developed a method based on impedances to estimate a comprehensive enough parameter search space. We also validated our assessment method with simulations of the linear system (in MATLAB) and nonlinear system (RT-lab).

In this work, we showed that step change disturbances due to static load variations do not complicate the stability assessment, and that we can drive the performance node states to their targets without needing to have the closed loop system be asymptotically stable. The placement tool was applied to the 13-node and 123-node IEEE feeders, as well as the PGE PL0001 feeder.[7]

Fig. 7 illustrates one type of result from the tool, applied to the IEEE 13-node feeder. In this function, the user sequentially places actuator-performance node pairs to understand how each additional placement informs the available locations for the next placement. In particular, at step k we choose a performance node location, which could be a known "sore spot" of the network where we want to mitigate overvoltage or undervoltage. Then a binary colormap produced with respect to the chosen performance node tells which nodes are "feasible" (colored green) and which are infeasible (colored red). In light of the color map info, the actuator node for step k is chosen. Then at step k+1 this process – choosing the performance node, generating the colormap, and choosing the actuator node – is repeated, until the user is satisfied with their accumulated configuration of actuators and performance nodes.

In some cases, the results are intuitive; in others, less so. The outcome of this effort is that we can answer the question of whether PBC with local PI controllers is feasible to implement on large, complicated circuits. The original approach described in the SOPO was simply to simulate many diverse PBC scenarios, including ones on feeders with many nodes. But because the number of possible scenarios grows combinatorially, and because it is not obvious when and where problems may arise, it would be unrealistic to make general claims based simply on a large number of successful simulations. Instead, this systematic approach can be applied to any candidate control scenario on any feeder, and predict its feasibility with confidence.

Figure 7: Sample results from PI Controller Placement Tool.

(a) In one of many simulation exercises, node 684 (square) is chosen as the performance node for which some phasor target is specified relative to the feeder head at the substation (node 651). The color map shows that all nodes on the feeder are feasible locations for siting an actuator: some might be more effective than others, but none create instabilities.

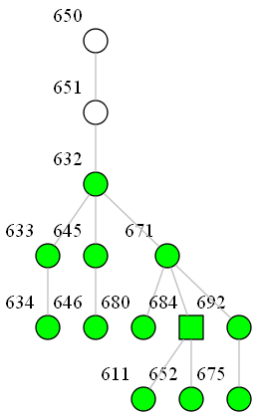

7 The methodology and some results are detailed in J. Swartz, B. Wais, E. Ratnam and A von Meier, "Visual Tool for Assessing Stability of DER Configurations on Three-Phase Radial Networks." Submitted to IEEE Powertech 2021. arXiv preprint available at arXiv:2011.07232

(b) We arbitrarily choose to place the first actuator at node 611 (gray), and, subject to that choice, examine other locations for actuators to jointly track node 684. Still, all nodes are feasible. We choose 652.

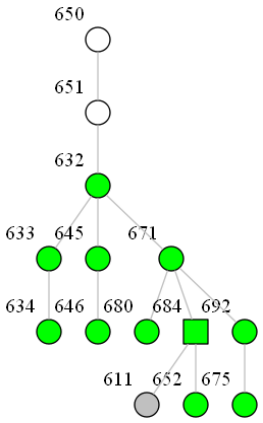

(c) With actuators at 611 and 652 (gray) acting on node 684, we introduce another performance node, 632 (square) and examine possible locations for actuators. Our options are now constrained. Node 680, 684, and 675 are infeasible (red), indicating that placing an actuator here could prevent the system from converging to the combined phasor targets at 632 and 684. We choose to place an actuator at 645.

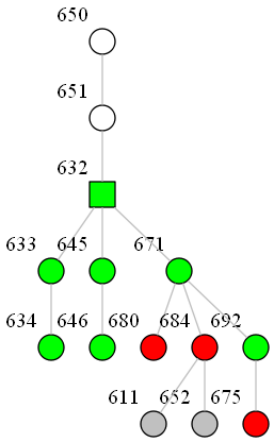

(d) With actuators at 611, 652, and 645 (gray), we examine actuator for tracking an additional performance node, track node 671 (square). At this stage, only node 671 (co-located) or 692 are feasible for placing another actuator that will not interfere with other control actions on the feeder.

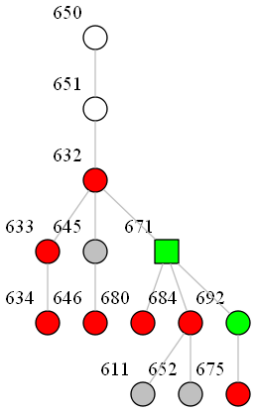

LQR Controller

The linear quadratic regulator (LQR) controller was designed to addresses several fundamental challenges of PBC: the coupling between the multiple input (P and Q injections on each phase) and multiple output (voltage magnitude and angle on each phase) or MIMO nature of unbalanced three-phase power flow; the nonlinear relationship between power injections and the local phasors; the inevitable model mismatch between the L-PBC's network model and the true model; and the effect that other loads and generators on the network have on the local voltage phasor.

In contrast to the PI and the RCAC controllers, the LQR uses a system model to create a linear, MIMO feedback policy. The LQR statespace includes both error and error-integrating states to account for model mismatch. The system model includes only the effect of the control input injection, treating all other load and generation on the network as an exogenous disturbance. In addition to the MIMO feedback policy, the LQR controller also implements an internal disturbance rejection loop, which estimates and counteracts the exogenous disturbance in real time. Fig. 8 illustrates how the LQR feedback, Disturbance Observer and network model interact.

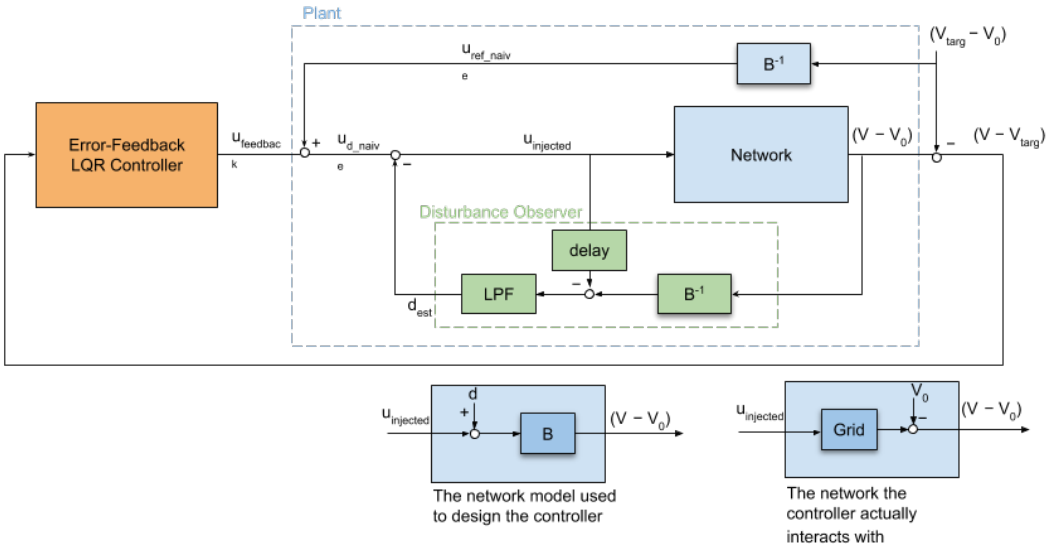

Figure 8: LQR controller schematic diagram

The purpose of the Disturbance Observer is to offset the effect of all of the exogenous power injections that are not controlled by the L-PBC. These exogenous power injections will occur regardless of the L-PBC's behavior. If the L-PBC does not take the exogenous injections into account, the L-PBC injections will not accurately track the phasor target, or rely too heavily on the LQR feedback, which is not as quick to adjust as an explicit disturbance observer.

The voltage disturbance is the effect of all of the other generators on the network on the L-PBC node's voltage. The voltage disturbance can be converted into an equivalent "effective power flow disturbance," which is the power flow that would create the voltage disturbance at the L-PBC node if the effective power flow went through the L-PBC node's effective impedance. The Disturbance Observer determines the effective power flow disturbance by subtracting the L-PBC-commanded power flow from the effective power flow (calculated by passing the measured voltage through an inverse power flow block), then passing the result through a low-pass filter (LPF) that serves to eliminate measurement noise from the L-PBC effective power estimate, as well as damping the dynamics of the inner loop. This damping is desirable, as the inner loop dynamics are not accounted for by the LQR error feedback.

From the Error-Feedback LQR Controller's perspective, the control output is u-uref and the state is V-V_{ref}. This allows the infinite horizon LQR to be designed with the standard LQR equations (including the canonical Discrete Algebraic Riccati Equation). Note the difference between the plant that is used to design the controller, and the plant that the controller interacts with when it

is actually deployed: the true network gives V for a control input u; however, the linearized power flow gives $V-V_0$. This discrepancy is resolved in practice by subtracting the substation voltage V_0 from the V that is measured/produced by the true network.

The Berkeley team also explored an alternative LQR control formulation. In this alternative formulation, the LQR controller uses an "endogenous" disturbance model, while the Disturbance Observer inner loop still treats all other load on the network as an exogenous disturbance. In this endogenous disturbance model, the expected value of the voltage phasor error state is expected to be equal to the value it had in the previous discrete timestep, rather than zero. The input u is the change in power injections, rather than the total power injections. The input matrix B is still the power flow linearization, but is now interpreted as a sensitivity matrix. We determined that this endogenous LQR formulation resulted in a faster, more effective feedback policy.

Furthermore, the endogenous disturbance model has a favorable interpretation in terms of the "optimality" of the feedback policy. The input cost matrix R now quadratically penalizes changes in the input, rather than the input itself. This provides the system operator with a knob that encourages the actuators to change slowly, which desirable for both hardware and system stability.

Both the LQR feedback policy and the Disturbance Observation require a network model. Distribution network models are often inaccurate or nonexistent. Furthermore, distribution networks can be reconfigured at any moment. Thus, the Berkeley team incorporated a Real-time Effective Impedance Estimator (REIE) into the LQR controller. This REIE allows the L-PBC to adapt to the time varying system online, using only measurements from the local PMU sensors.

The REIE can be thought of as a parallel system on the L-PBC that produces an estimate of the effective impedance between the L-PBC sensor node and the substation (assumed to behave as an infinite bus) from the local voltage measurements and power commands. This effective impedance estimate is maintained using exponentially-weighted recursion, and therefore provides an estimate of the effective impedance in real time, without requiring a batch of phasor measurements to be stored locally.

The voltage at the L-PBC node is determined by all of the power flows on the network. Thus, the effective impedance cannot be estimated by simply dividing the local voltage by the local current injection. Instead, the small signal impedance is used, calculated using the change in local voltage and change in current injection at the L-PBC node between the present and preceding timestep. While the loads on the rest of the network will still affect the change in voltage, the hypothesis behind the REIE method is that the effect of these loads will be zero-mean and uncorrelated, and can therefore be effectively eliminated via time-averaging regression to produce a stable and unbiased estimate of the effective impedance.

The performance of the LQR controller was successfully validated in HIL testing.

RCAC

In a parallel effort, the University of Michigan developed and refined its Retrospective Cost Adaptive Control (RCAC) algorithm.

Analogous to the effective impedance estimator described above, 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. The controller does not rely

on any externally provided information about impedances within the network. (Note that impedances given by the feeder model are used to produce the behavior of the plant in simulation, but this information is not used by the controller and is not required in a physical implementation.)

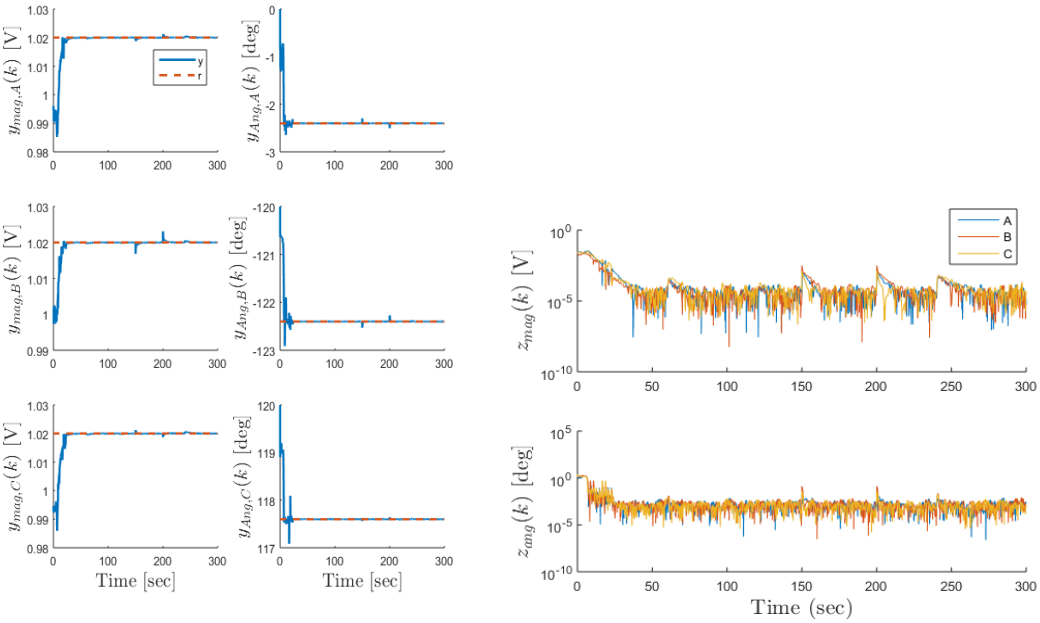

Figure 9: Phasor tracking by RCAC controller.

The RCAC algorithm quickly proved capable of performing PBC. Fig. 9 illustrates sample results for disturbance rejection of voltage magnitude and angle on three phases on the IEEE 13-node feeder, with the right side showing the phasor tracking error on a logarithmic scale. While the initial settling time when approaching the phasor target from an uncontrolled state is 10-15 seconds, subsequent actuation is much faster.

Refinements then focused on simplifying the RCAC structure and modeling information. Typically, a feeder would be impulsed several times in order to obtain sufficient modeling information. One unexpected finding for the IEEE 13-node feeder was that that an approximation of the RCAC-specific target model was sufficiently robust to changes in actuation and performance nodes, removing the need to impulse the system repeatedly for different control scenarios.

Like the other L-PBC approaches, RCAC was demonstrated across multiple test cases with the goal of challenging the controller. For example, RCAC was tested in cases where the actuation and performance nodes were close, or electrically distant; it rejected both step disturbances and sinusoidally varying distrubances at multiple, electrically separated nodes. The Michigan team also demonstrated the feasibility of a decentralized architecture, where two RCAC subcontrollers

follow setpoint commands at electrically separate and distinct performance nodes without adverse interaction.[8]

An essential part of RCAC is the calculation of a recursive least squares (RLS) solution for the adaptive gain update. Implementations of classical RLS on time-varying systems have trouble, however, since the algorithm equally weighs all previous data points, when the truth is that some data contain more information about the current system characteristics than others. This makes it impossible for the system to adapt to changes after a certain amount of time. Historically, this led to the introduction of the "forgetting factor," a constant between 0 and 1, which is used to forget previous data by treating it as though it were exponentially less informative, the further back in time it is located. Constant forgetting factors, however, can sometimes lead to numerical divergences, especially when the controller inputs lack persistency – that is, if they lack enough information to properly ascertain certain system characteristics from recent data.

An important refinement to RCAC focused on developing variable forgetting factors for RLS, which can both change the rate of forgetting and the directions in parameter space in which forgetting is done. This research has led to success in overcoming issues with loss of persistency, as well as RLS implementations that can quickly adapt to sudden changes in system characteristics – rather than adapting only slowly varying changes, as is the case with constant forgetting factors.

These general results have important implications for the robustness of PBC. For example, variable-rate forgetting may help address catastrophic situations or significant contingencies, such as the failure of actuation nodes or sudden network topology changes unknown to the supervisory controller. Variable-direction forgetting might be able to overcome situations where performance node signals, either gradually or suddenly, become unavailable.

There are many more hypothetical situations than could be simulated within the scope of this project. However, the findings by the Michigan team give us reason to be confident that, with its inherent adaptability, the RCAC approach is broadly applicable to PBC and could be readily implemented in practice.

Phase Coupling Challenge

In the process of designing a particularly challenging but realistic test case to compare the performance of different controllers, the team discovered a model conversion error that proved instructive about the vulnerability of PBC to extreme phase coupling.

In a comparison simulation on the IEEE 123-node feeder, actuation commands are sent to node 49 to track a phasor target at node 49, a simple configuration. A 25-kW, 330-kW, and 150-kW square-wave disturbance occur at node 49 on phase A, B, and C respectively, from 800 timesteps to 1300. We found the genetic algorithm tuning method for the PI controller resulted in an unstable system. When we manually tuned the controller gains we were able to stabilize the system, but unable to drive all three phases to their phasor targets. Even more surprising, the RCAC controller resulted in similarly poor performance.

8 The RCAC approach and early simulation results are presented in A. Ul Islam, E. Ratnam and D. Bernstein, "Phasor-Based Adaptive Control of a Test-Feeder Distribution Network." IEEE Transactions on Control Systems, 2019.

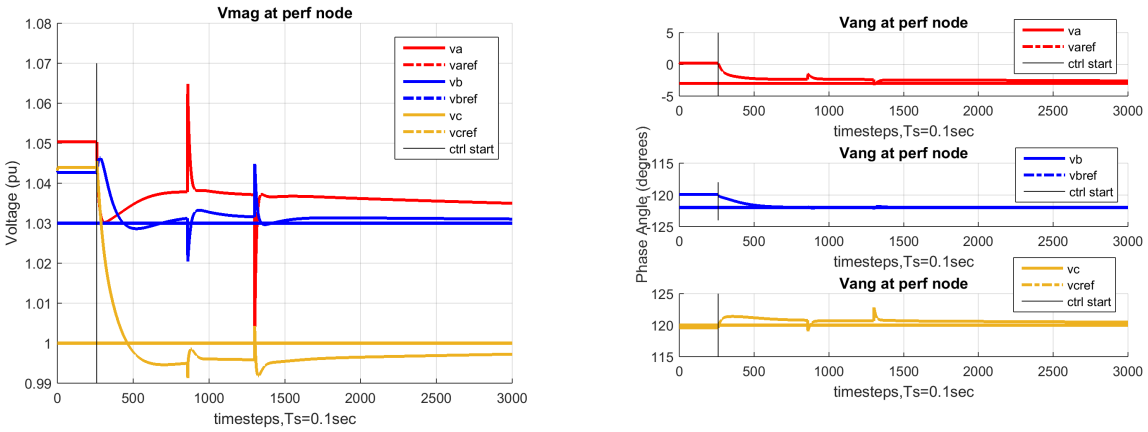

Figure 10: Voltage magnitude (left) and phase angle (right) at node 49a/b/c on incorrect 123NF when using PI controllers to track 3-phase phasor targets. Phase coupling due to unrealistically high mutual impedance prevents convergence.

We then discovered a phase translation error in the 123-node feeder that created an unrealistically high mutual impedance value for one circuit branch, meaning that voltages on one phase would be highly sensitive to power injections on another. After correcting this error, it was easy to tune both the PI and RCAC controllers to track the phasor on this configuration. This change from a problematic to a well-behaved model demonstrated that phase coupling can be a direct obstacle to PBC tracking. Though it seems unlikely that a real feeder would exhibit behavior quite as pathological as the incorrect model, it is worth noting that the vulnerability exists and equally affects different L-PBC controller algorithms.

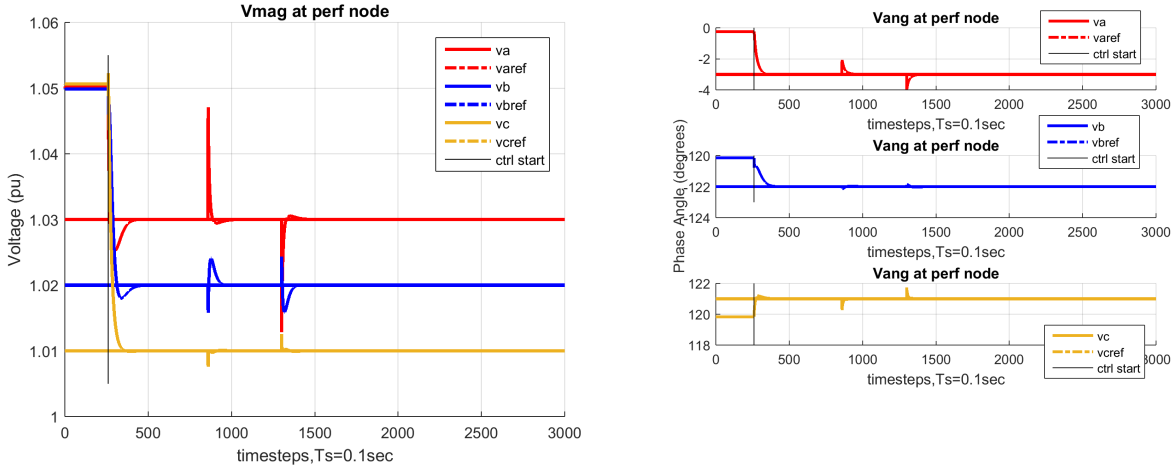

Figure 11: Voltage magnitude (left) and phase angle (right) at node 49a/b/c on corrected 123NF when using PI controllers to track 3-phase phasor targets.

Stability Investigation

The team also completed a separate, general investigation of PBC controller stability. The danger of instability is particularly pertinent for the PI local controllers, which do not incorporate network information, and make independent single-input, single-output (SISO) decisions with separate treatment of variable pairs. To address this problem systematically, the Berkeley team investigated a multitude of parameter variations to empirically determine the stability properties of the system. This was done with simulation runs in Python, where OpenDSS was used to run the three-phase network power flow in each simulation step. The simulations were conducted on simple 2 and 4 node networks to develop intuition, and then on the 13-node IEEE feeder.

Because the PI controller makes decisions in a SISO manner, any source of coupling between the actuators and the state variables that are not assigned to that actuation are a potential source of instability. The goal of the empirical testing was to isolate various forms of coupling that might lead to instability during phasor-based control using a PI controller. The Berkeley team identifying four different types of coupling/potential sources of instability:

1. P-Q Coupling – Reactive power (Q) affects both voltage magnitude and angle, not just magnitude. Real power (P) affects both as well, not just angle. The approximation that Q determines voltage magnitude and P voltage angle works best on reactive networks, and may fail on resistive networks. Since our PI controllers use this assumption, it could lead to instability when it proves untrue. This type of coupling should thus be related to high R/X ratios.
2. Phase Coupling – Although the three phases of a network are approximated to be electrically distinct, they are in fact connected. When one phase on a network falls out of balance, the effective impedance of the network increases and this affects the other phases as well. This could lead to coupling back and forth between actuation and sensing on different phases. This type of coupling should be related to the cross-impedance of the network.
3. Multiple Actuator Coupling – When multiple actuators are responding to a single performance node, they may interfere with one another.
4. Multiple Performance Node Coupling – When multiple performance nodes are attempting to meet their targets on the same circuit, they may interfere with one another.

The simulations demonstrated that the primary factors contributing to instability in the network were the P-Q coupling from large R/X ratios, and phase coupling from the distribution line cross-impedances. Furthermore it was determined that the two types of coupling were related—i.e. R/X ratios and network coupling amounts that were independently benign can create instabilities when combined.

With added complexity, instability can appear in a delayed or temporary fashion. In a single-phase case, instability always appeared clearly from the beginning of the simulation. With the addition of interaction between phases, instability appears at much lower gains. It can also appear after a delay, and form an oscillating pattern, as shown in Fig. 12.

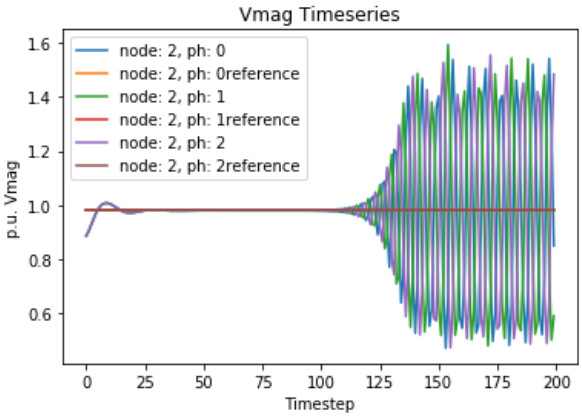

Figure 12: Simulation on two-bus system showing control instability due to coupling with interaction between phases. Because the phases are balanced at first, the instability takes 100 timesteps to emerge. Then, by timestep 150, the instability has settled into an oscillating pattern.

When actuation and sensing are separated, or multiple performance nodes are introduced, instability may take even stranger forms. For example, with sensing at node 675 and actuation at node 634, instability seemed to emerge and retreat multiple times throughout the simulation.

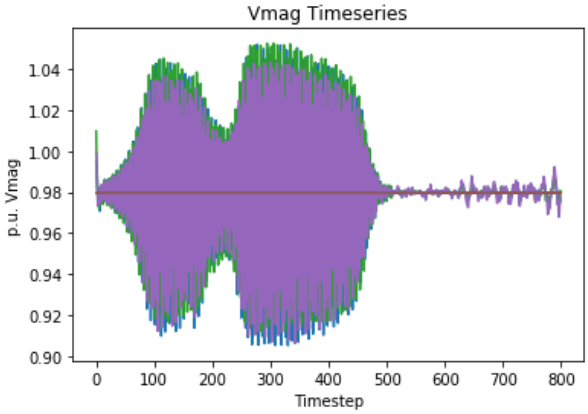

Figure 13: Simulation with actuation placed far from performance node on the IEEE-13 feeder. Oscillations are observed, followed by a period where stability is regained.

Instability tests were conducted with the PI and LQR controllers. No instabilities were observed with the LQR controller with any realistic impedance values. This can be attributed to the fact the MIMO feedback of the LQR controller inherently takes into account the network's R/X ratio and the network coupling.

It is worth noting that adding actuation nodes that track the same voltage phasor did not seem to adversely affect stability in our simulations. On the contrary, it was common to see slight increases in stability as the total actuation was distributed throughout the feeder rather than coming from a single node.

Subtask 6.2.4: Handoff L-PBC (M19-M21): Prepare L-PBC algorithms for implementation on Controller Platform, to produce L-PBC output signal based on simulated supervisory control signal specifying phasor target.

The original hypothesis was that PBC would be implemented as one of many functionalities on an existing commercial DER controller. Instead, the project developed its own DEGC platform, detailed below. The L-PBC optimization was run by the Berkeley team on a standard fitPC, with no handoff required.

Subtask 9.1.5: Test PBC (M25-M27): Test circuit models and associated phasor targets for PBC applications, prepare models for integration within the OPAL-RT simulator.

Circuit models were converted into ePHASORsim throughout the project. Once converted, integration with the OPAL-RT simulator was straightforward.

During the last two quarters, after the task of preparing models for simulation was complete, OPAL-RT also helped troubleshoot a convergence problem that sometimes occurred on one of the PG&E Taxonomy feeders. It was suspected that the problem resulted from mutual coupling between phases, where, for example, controller actuation on Phase A would significantly affect voltages on Phases B and C.

To diagnose the issue, OPAL-RT and UCB performed a systematic characterization of L-PBC controller gain for all 341 nodes on the PL0001 PG&E distribution feeder to study the mutual coupling effect, using co-located performance and actuator nodes. For the test, reactive power was injected on a particular phase on a node and the effects on the voltage magnitude and angles on other phases of the same node were analyzed. Based on the acceptance criteria needed for the L-PBC controller characterization, the nodes were then flagged as usable or not usable for co-located L-PBC control.

The finding that phase coupling could prevent PBC from working effectively on an unbalanced feeder had not been anticipated. However, the phenomenon was successfully diagnosed and led to the ability to prioritize PBC nodes on a feeder with high mutual impedances across phases and predict which, if any, are unsuitable for siting actuators.

Subtask 9.1.6 Refine S-PBC (M25-M27): Refine S-PBC algorithms for robustness in the presence of missing and/or corrupted data (e.g., compute feeder-specific tolerances for PBC targets so that PBC algorithms operate within known grid constraints), produce supervisory control signal and associated tolerances in response to input data. Refined S-PBC algorithms to be evaluated on desktop simulations outside of the OPAL-RT simulation environment.

PBC was shown robust with respect to missing and bad data inputs. Specfically, we demonstrated smooth recovery when the local controller finds that it is infeasible to attain the phasor target – which could result from a poorly calculated target by the supervisory controller, from a phasor measurement discrepancy, or from a discrepancy between expected and actual available actuation resource, causing the controller to saturate. In such a case, the local controller sent an “I Can't Do It” (ICDI) signal to the supervisory controller. S-PBC then responds with an

adjusted phasor target. This interplay was successfully tested with both controller-in-the-loop (CIL) and HIL environments. Recurring or persistent ICDI signals would allow the supervisory controller to flag potentially corrupt data.

Subtask 9.2.5: Refine L-PBC (M25-M27): Refine and improve L-PBC control algorithms for robustness in the presence of missing and/or corrupted data. Test performance of L-PBC algorithms when increasing numbers the number of nodes (compare the performance of L-PBC when controlling 10 nodes against controlling 1000 nodes) on interconnected distribution feeders. Refined L-PBC algorithms to be evaluated on desktop simulations outside of the OPAL-RT simulation environment.

ICDI

Robustness of the local controller with respect to missing or corrupt data was demonstrated using the PI controller, in both controller-in-the-loop (CIL) and HIL environments. In the event of missing data such as a disconnected μPMU source, the local controller simply fails safe to allow default operation of the distributed resource.

Fig. 14 illustrates an example of a successful phasor target revision by S-PBC in repsonse to an ICDI signal from L-PBC. In this example, a battery resource is dispatched to increase its power consumption (negative injection), charging the battery to draw the nodal voltage down to a low value of 0.92 p.u. (Note that, like many of our simulation scenarios, this case was not designed to be realistic, but to challenge the controller and stress the actuators to their limits.) The phasor target is suddenly changed from its initial value (yellow), which was being tracked by the PI controller in the face of large disturbances, to an unreasonably low value (green) that cannot be achieved even with the battery at its maximum charging rate. Upon receiving the ICDI signal from L-PBC, the supervisory controller issues a revised target (red), on which the controller smoothly converges. This HIL test result closely matches the previous CIL simulation.

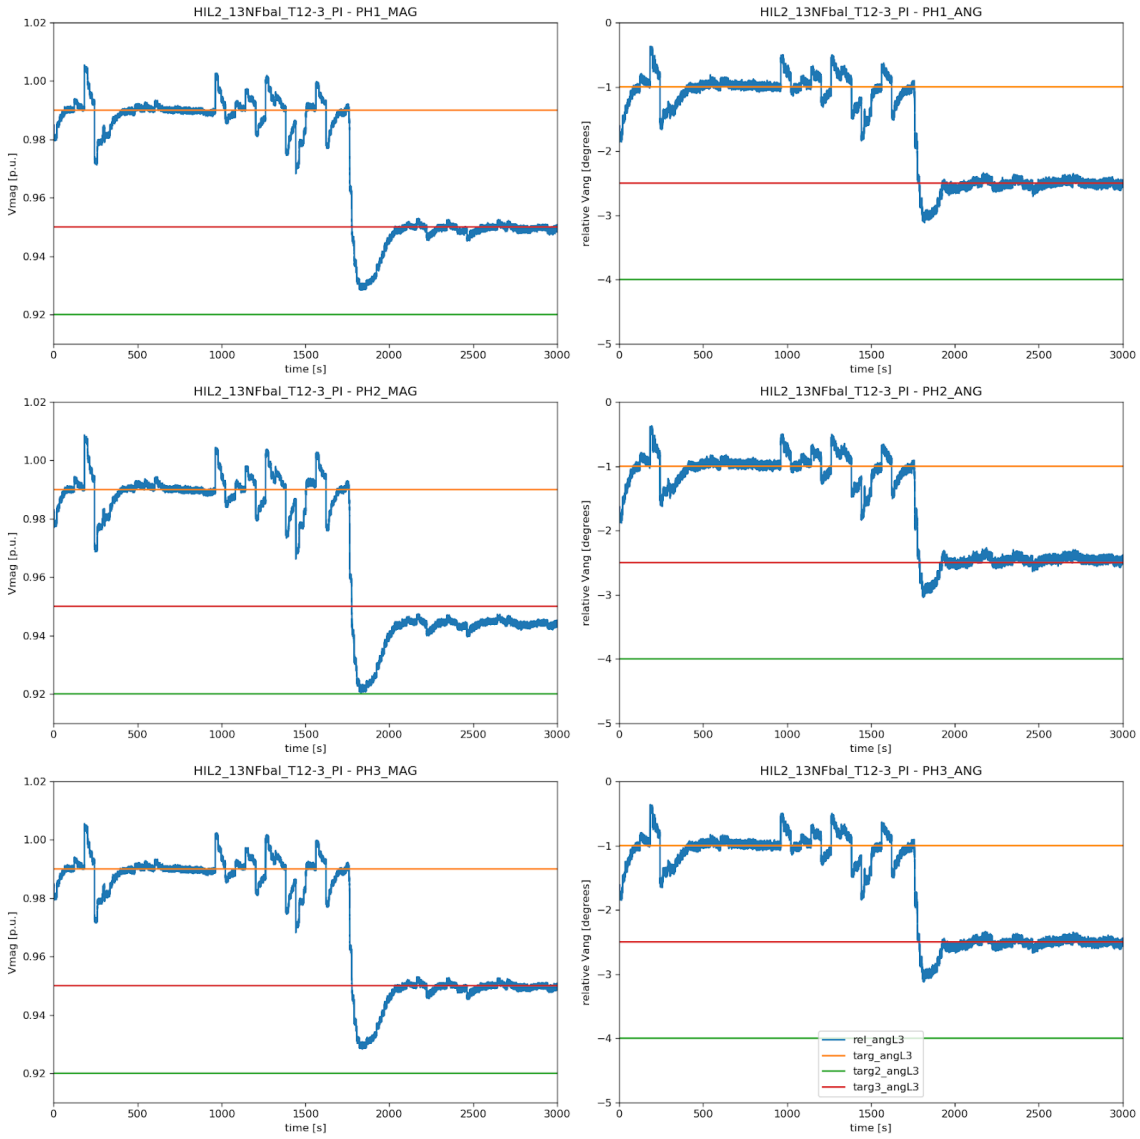

Figure 14: Test scenario 12-3 on the 13-node balanced feeder, showing the controller recovering from the "I Can't Do It" condition with an unrealistic phasor target.

Scaling Controller Nodes

The team had hoped to perform simulations on an 8500-node circuit prepared by OPAL-RT, but ran out of time to do this. Weighing the value of experimenting on a large circuit absent any specific hypotheses of how this circuit would behave differently than those already tested (and at the risk that this new circuit could take considerable time to debug) against a specific comparison of single- versus multi-node control, the project team decided that its effort would be better spent on the latter. While we did not attempt to control 1000 nodes, we carefully compared the performance of controlling 10 distinct nodes against controlling a single node, on one of the larger and most challenging circuits.

Two types of complex scenarios were simulated on the 341-node PG&E feeder. To define the configuration of actuators and performance nodes for each scenarios, we first selected ten

approximately evenly distributed node locations throughout the feeder. For the first scenario, we assigned a single performance node in the center of the feeder at node 300020414, and employed a separate two- or three-phase actuator at each actuator node to collectively track the same phasor target. For the second scenario, we assign every actuator node to also be a performance node, and employ every actuator to track a phasor target at its own node. In both simulations we include 100% PV penetration at every load node (114 nodes).

Node number	Phases present	Distance from feeder head (|Z| in ohms, averaged across phases)
300033983	A/B/C	1.276
300063918	A/B/C	3.06
300062294	A/B/C	2.863
300233573	A/B/C	3.592
300062321	B/C	4.08
300020632	A/B/C	4.535
300062066	A/C	3.695
300004134	A/B	4.394
300062296	A/B/C	5.354
300053281	A/B/C	5.206

Table 2: 10 nodes on PL001 selected for Scenarios 1 and 2

We ran the PBC configuration feasibility tool on both scenarios, verifying that there exist controller gains that will ensure all phasor targets are met. The tool determined that while the first scenario is not very sensitive to where the actuators are located, the feasibility of the second scenario depends on adequate spacing between performance nodes, which prevents co-located pairs from destabilizing each other.

We then ran both scenarios in RT-LAB with ePHASORsim. For Scenario 1 (multiple actuators, single performance node), we select a plausible three-phase phasor target of [0.97 0.98 0.99] per-unit voltage magnitude and [-2 -122 122] degrees, apply actuator limits of 400kW, and include the time-varying load/generation data. We tune the controller gains using the genetic algorithm auto-tuning procedure, which makes use of sensitivity values computed from step change response data.

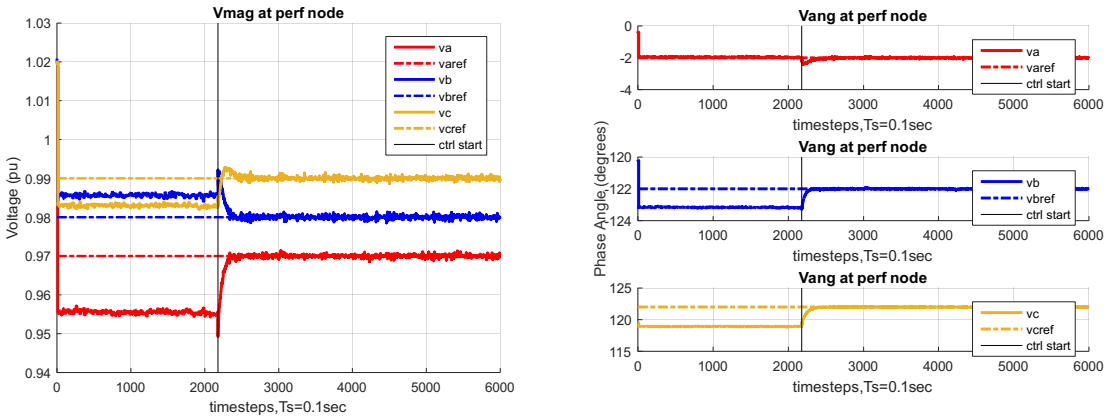

Figure 15: Voltage magnitude (left) and phase angle (right) at the performance node

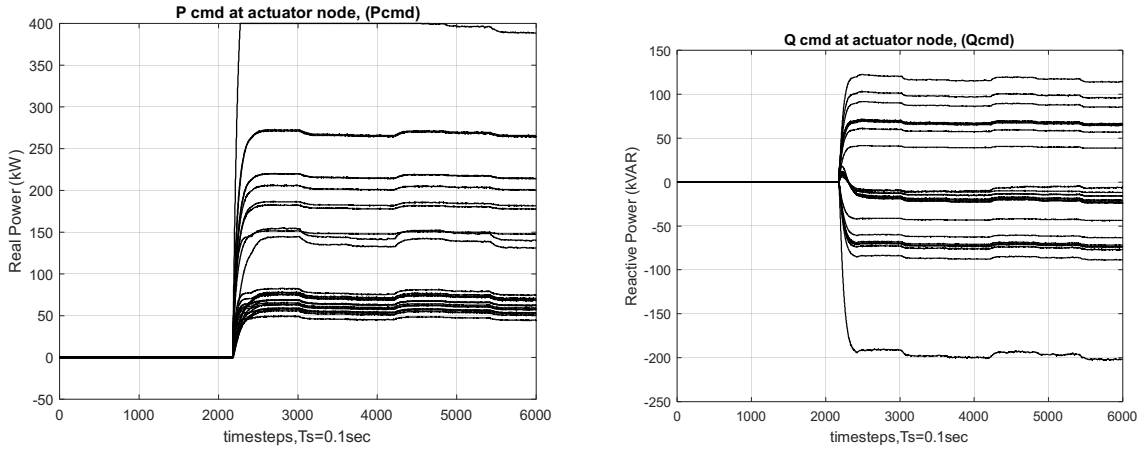

Figure 16: Real (left) and reactive (right) power commands computed for all ten actuator nodes. The real power on Phase C of node 300033983 saturates.

Scenario 1: When controllers turn on at 210 seconds, the actuators each provide real and reactive power, causing the voltage phasor to settle to its target within 15 seconds with minimal overshoot and oscillations. We observe that when the actuator on phase C of node 300033983 saturates (reaches 400kW), the other actuators continue to increase their actuation so that the phasor target is still met. We are able to tune this multiple actuator configuration without much difficulty. Specifically, a range of controller gain sets found by the tuning algorithm result in smooth convergence to the phasor target. This observation supports earlier findings in smaller-scale simulations that multiple actuator coupling when tracking a single performance node is not a significant problem.

Scenario 2: This test with multiple independent performance nodes highlights the importance of the S-PBC optimization, because with ten phasor targets, it is critical for the them to be computed in a consistent manner – that is, without internal contradiction as to the physical feasibility of power flows between nodes. In this test case, the targets are computed optimally with respect to the objective of phase-balancing. Between 280 and 390 seconds, we introduce an additional challenge to the controllers by creating a square wave disturbance, corresponding to an aggressive 50% increase and decrease in load, at every performance node simultaneously. As in the previous scenario, we enforce 400 kVA actuation limits.

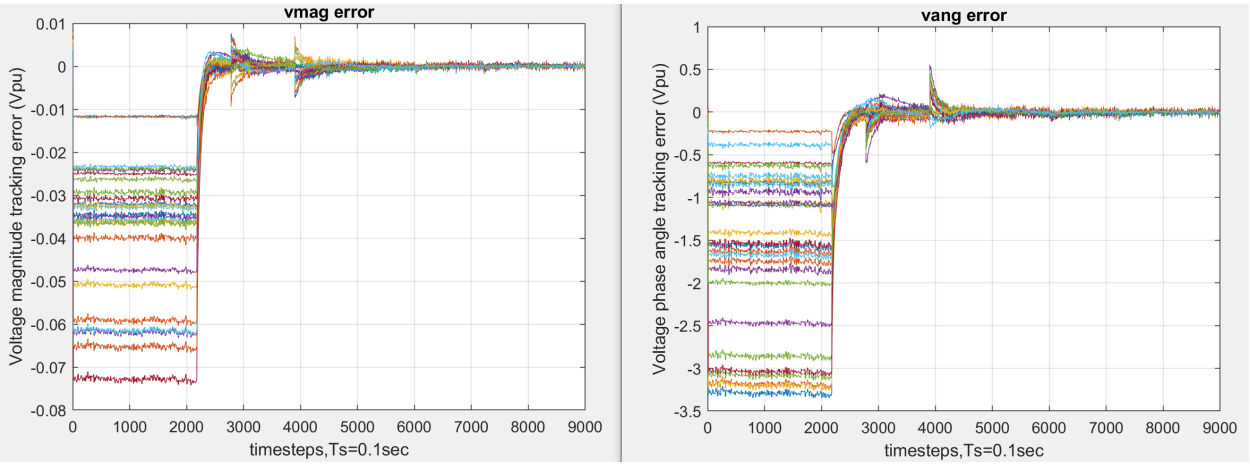

Figure 17: Voltage magnitude tracking error (left) and voltage phase angle tracking error (right) for all 10 performance nodes.

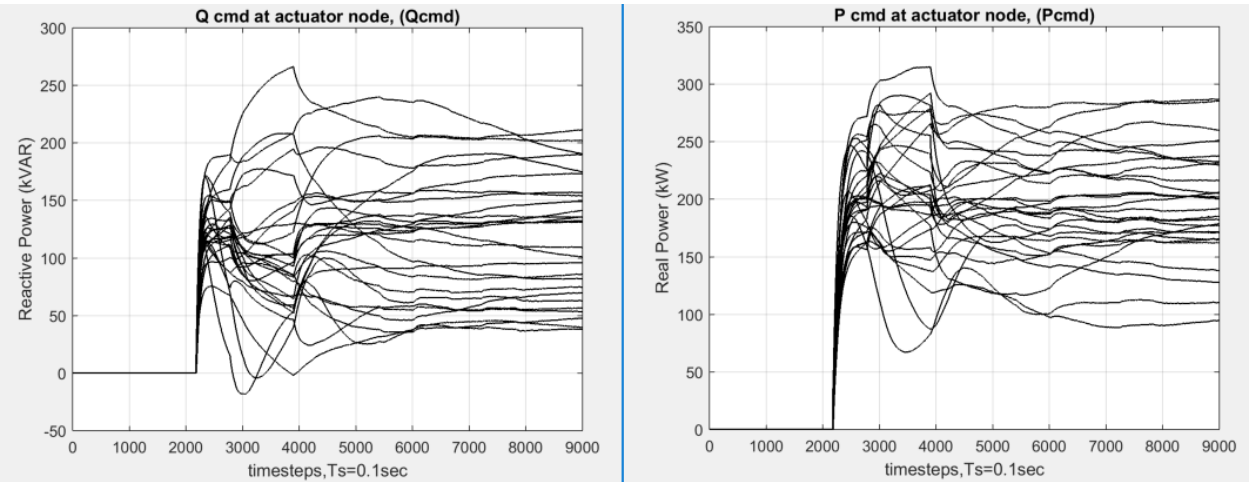

Figure 18: Reactive (left) and real (right) power commands computed for all ten actuator nodes.

When the controllers turn on at 210 seconds, it takes longer than in the previous scenario, about 100 seconds, for all the voltages to settle to their phase targets. Partly this is due to the ten different phasor targets, and partly due to the large square wave disturbance introduced before the system settles. It was harder to tune controller gains for this scenario than the previous. Even after the genetic algorithm determined the appropriate relative strength of controller gains between actuators and accounted for coupling of real and reactive power, all gains were reduced by a percentage to achieve successful (yet slower) convergence.

Based on the success of showing PBC in this complex yet also realistic scenario, the team chose it for the OPAL-RT demo, discussed below.

Controller Implementation and Data Infrastructure

Subtask 3.1: Compatibility Assurance (M1-M12): Coordinate and confirm that L-PBC algorithm requirements and platform capabilities are compatible. Specifically, verify that expectations regarding data formats, communication protocols, data rates, communication latencies and computation latencies for PBC are consistent.

The original plan was to deploy PBC as one of several voltage control algorithms on the Doosan GridTech Intellicant Controller (DG-IC). The controller had been developed to communicate with a single μPMU using the IEEE C37.118 protocol and would send inverter commands over Modbus. In the first project year, Doosan extended the DG-IC capability to support communication from at least two μPMUs simultaneously. Because some of the documentation is proprietary, and also because this implementation strategy was not ultimately pursued after Doosan's departure from the project, we do not further detail it in this public report.

An important takeaway from this effort, however, was that deploying PBC within the context of an existing commercial DER controller posed no fundamental challenges that would cast doubt on the practical viability of the paradigm.

Subtask 7.2: Specifications (M13-M18): Develop testing protocol and performance criteria for local controller.

Subtask 7.3: Design controller (M13-M22): Design local controller to interface with specific devices (e.g., smart inverters and battery storage systems), and prepare to integrate local controller for HIL testing.

Instead of integrating the Doosan GridTech controller with the HIL environment, the team created a simple, adaptable design using minimal, standard resources. One deployment option for the L-PBC is to run on a Raspberry Pi; another is to run the controller code on a fitPC. In collaboration with LBNL, the Berkeley team successfully prepared the hardware and software for deployment at the FLEXLAB.

Subtask 7.4: Implement S-PBC (M18-M24): Implement S-PBC algorithms on a virtual optimization platform capable of interfacing with OPAL-RT's simulator and a physical controller, to produce a supervisory control signal in response to simulated input data.

Subtask 7.5: Implement L-PBC (M18-M24): Implement L-PBC algorithms on a physical local controller, producing L-PBC output signal in suitable format to be sent to physical devices, based on simulated supervisory control signal specifying phasor target.

The Berkeley team took over responsibility for developing the controller platform after Doosan opted out of the project at the end of Year 1.

We built a data infrastructure for the HIL testing environment that was both practical for integrating with the FLEXLAB setup, and promises adaptability and scalability to larger implementations in the field. To this end, we were able to leverage previous research and development efforts by UC Berkeley.

A key element is the publish-subscribe WAVEMQ message platform, which implements the Wide Area Verified Exchange data bus (WAVE) for decentralized authentication and authorization.[9] WAVE implements a Merkle tree, and can be thought of as a less cumbersome version of Blockchain. This had been previously applied in the context of an eXtensible Building Operating System (XBOS) data infrastructure, where the various devices to be addressed would typically be diverse loads in a large commercial building. It is adapted here to the specific data requirements associated with high-resolution streaming phasor measurements from micro-PMUs.

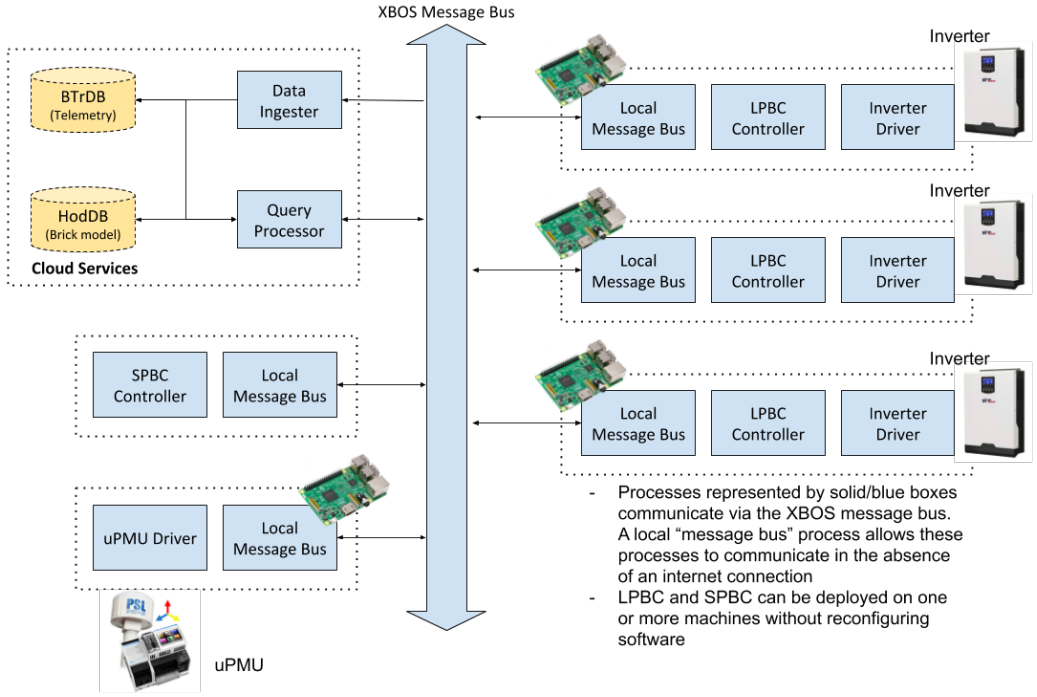

Figure 19: Communication infrastructure for HIL testing.

Of course, the deployment also leveraged micro-Phasor Measurement Units (μPMUs) developed under the 2012 OPEN ARPA-E funded project[10] led by Berkeley/CIEE, as well at the Berkeley Tree Database (BTrDB) developed in the same project. BTrDB has been commercialized by

9 See M. P. Andersen, S. Kumar, M. AbdelBaky, G. Fierro, J. Kolb, H.-S. Kim, D. E. Culler, and R. A. Popa, "WAVE: A decentralized authorization framework with transitive delegation," in 28th USENIX Security Symposium (USENIX Security 19). Santa Clara, CA: USENIX Association, Aug. 2019, pp. 1375–1392. [Online]. Available: https: //www.usenix.org/conference/usenixsecurity19/presentation/andersen

10 Award DE-AR 0000340.

project partner PingThings, and it underlies their PredictiveGrid™ platform.[11] This is indicated as "Cloud Services" in Fig. 19, which could of course include other platforms and applications.

We generalized the infrastructure built for HIL testing into a secure, distributed control platform titled the open-source Distributed Extensible Grid Control (DEGC) platform, which can facilitate arbitrary distributed control implementations, including PBC.[12]

DEGC provides an extensible software and communications platform for DER control that facilitates distributed computation, as well as communication between and among sensors and controllers. The DEGC platform consists of distributed assets implemented as DEGC processes, and a secure, high-reliability publish-subscribe message bus that facilitates communication between the DEGC processes. Both the publish-subscribe message bus and the abstractions of the DEGC software platform are designed to provide an extensible control platform to control system designers/engineers. Extensibility is a key property for DER control methods. Specifically, the extensibility provided by the DEGC publish-subscribe message bus allows the DEGC PBC implementation to easily add and remove controller nodes and actuators.

Fig. 20 describes more abstractly how PBC was implemented using the DEGC platform. The DEGC platform implements every actor as a DEGC process. Within the DEGC process, there are four different classes: Controllers, Remote Actuators, Data Sources, and Data Sinks. The S-PBC controller, and each L-PBC controller, are implemented as Controller classes. Each of the participating controllers/sensors/actuators communicates with each other over the publish-subscribe message bus, via the protocol that is defined in the DEGC process layer. This framework is made available as open source code.

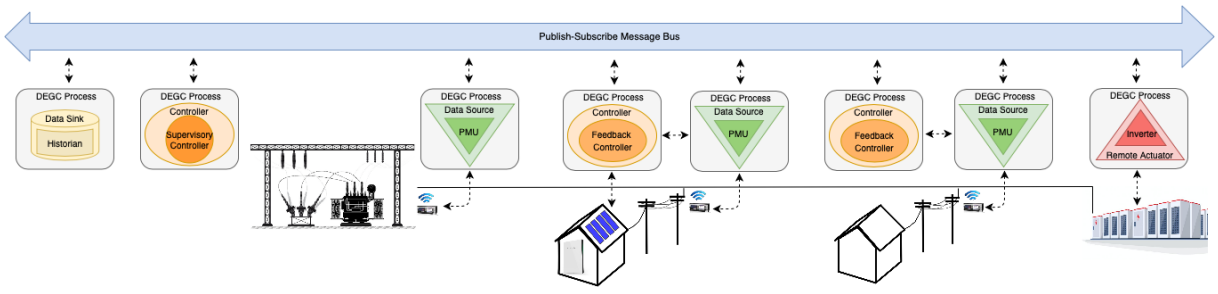

Figure 20: Generalized DEGC Implementation for PBC.

11 see https://www.pingthings.io/

12 DEGC is presented in G. Fierro, K. Moffat, J. Pakshong and A. von Meier, "An Extensible Software and Communication Platform for Distributed Energy Resource Management." IEEE SmartGridComm'20, Nov 11-13 2020.

Subtask 10.5: Commercialization Plan (M28-M30): Develop commercialization plan for local controller including licensing arrangements with partners, new partnerships, private or public fundraising, customer demonstrations, spin-off by core team members, etc.

While DEGC was developed specifically as an enabling technology and used for hardware-in-the-loop testing of Phasor-Based Control, it is in fact agnostic to control paradigms. Because of its broad applicability across Distributed Energy Resource (DER) devices, control strategies and algorithms, we believe that DEGC holds considerable promise for near-term commercialization – likely sooner than PBC itself. Therefore, the team opted to explore commercialization pathways for DEGC as a platform.

DEGC Platform Features

DEGC is built from the ground up to provide the DER owners with DER control security and energy use data privacy. Rather than relying on a central authority security, DEGC implements a distributed, zero trust security model. Resolving the security questions at the platforms foundation drastically reduces implementation complexity once the system reaches massive adoption.

The DEGC platform implements a publish-subscribe (PubSub) message bus which provides an extensible, asynchronous communication channel between the demand response or virtual power plant operators and all of the DER. The PubSub message bus is built on top of a decentralized WAVE security framework. Recently produced by researchers at UC Berkeley, and unique to the DEGC platform, WAVE provides three salient security features that make it well suited for privately-owned DER control:

1. Decentralized security
2. Graph Based Authorization (GBA)
3. Partial/granular permissions and privacy

Each of these features make it easier for the DEGC platform to keep the clients' DER secure, and their data private.

The decentralized security model maintains privacy and autonomy by default. So when a DER is connected to a demand response or virtual power plant operator, no permissions are given and no data access is immediately conferred. The DER owner can choose to give the demand response or virtual power plant operator certain control permissions, based on the incentives of the demand response or virtual power plant program. The DER owner can revoke these permissions at any time.

The GBA scheme makes it easy to reconfigure permissions once they are given. For example, take the case of a DER for which the owner has delegated control permission to a DER aggregation service, which has subsequently delegated permissions to a demand response program as well as a virtual power plant. If the DER owner would like to exit all of the programs, the owner only needs to remove the permission that it gave the aggregation service. All subsequent permissions that were given are removed accordingly. This makes permission management straightforward, even if the market for DER permission is complex.

Partial/granular permissions allow a DER owner to grant specific permission to only what the DER owner wants to grant permission to. For example, this allows the DER owner to grant permission to control the device but not permission to access the device's operational data.

Individual DER devices at a given location (e.g. a participant's house, office building, building complex, etc.) are connected to the DEGC hub via a LAN network (e.g. Wifi, Zigbee, or modbus), or directly to the Internet. The DEGC hub connects to the demand response or virtual power plant operator via a WAN network (the Internet). Thus, the DEGC platform provides a secure layer of separation between DER devices and the Internet, which is a significant cybersecurity hazard. Additionally, the DEGC platform provides a stopgap for denial of service (DOS) attacks that can overwhelm the DER devices with limited computational power.

DEGC Market

As readers of this document are well aware, the power distribution landscape is changing. High penetrations of distributed solar PV along with large-scale electric vehicle (EV) adoption present new levels of utilization and strain on distribution networks. These new strains can manifest as voltage or line flow constraint violations, which in turn pose serious threats to safety and service reliability for electric utility customers.

Preventing such violations will require one or more of the following strategies: (1) Placing conservative limits on distributed solar and EV adoption, as expressed in maximum feeder hosting capacities; (2) "Gold plating" distribution networks by upgrading infrastructure to accommodate maximum EV load or maximum PV generation at minimum load, even if these scenarios occur only a few times a year, if at all; or (3) Installing a communication and control platform that can actively control DER including PV, EV chargers, battery energy storage and other controllable loads in order to avoid constraint violations.

Climate goals and existing policies at the State level – along with Federal policy such as FERC 2222 – will rule out Option (1) in the long run. Regarding Option (2), upgrading distribution network equipment is inevitable, but a full "gold plating" of the grid to satisfy the worst case scenarios would waste ratepayer money. Option (3), combined with strategic distribution network upgrades is the most cost-efficient way of dealing with the intermittent nature of distributed solar and mass EV adoption. Active DER control schemes, including demand response and virtual power plants, coordinate DER to provide energy and/or reduce stress on the distribution grid when called upon.

The vast majority of DER are privately owned, including the growing set of controllable loads that constitute the "Internet of Things." Unfortunately, a barrier exists between privately owned DER and active control participation. The majority DER owners do not have the time to manually participate in active control programs that require them to make energy use decisions on minute-wise or evenly hour-wise timescales, and second-wise decisions are out of the question. For example, Southern California Edison, as well as a number of other utilities in the US, has a text message-based demand response program. The text message-based demand response programs have obvious limitations for speed and scalability, though, because they require human engagement. In order for privately owned DER to participate in an active control scheme on a mass scale, the control decisions must be automated. However, existing frameworks for automation still pose integration challenges. For example, even the OpenADR framework, intended to promote interoperability, depends on incorporating suitable communication and control capabilities within product hardware in order to allow participation within this

ecosystem. There remains a need to coordinate across silos with numerous types of devices and control strategies.

While many "smart home provider" (SHP) companies, such as Nest, have focused on this market, they have not largely been successful in automating DER at scale. This is for several reasons: first, few products actually have integration with the information required to make control decisions for grid stability, such as an OpenADR signal or realtime pricing information. Second, unless a user has invested in products from the same vendor, coordinated control of these resources is hard to achieve. Although some recent industrial efforts, such as Connected Home Over IP (CHIP), attempt to address this interoperability problem, these efforts concentrate on smart home appliances and do not provide control over general DER resources. There remains a need to integrate across load, generation and storage devices over a range of scales.

In addition to interoperability across DER and smart device platforms, in order for privately owned DER to participate in an active control scheme on a mass scale, the DER owner must choose to participate. Thus, the DER coordination must be designed with the DER owner in mind. We have identified several specific needs of DER owners that are difficult to satisfy simultaneously:

1. The DER owner must be kept in the loop. That is, the resource cannot be turned on or off without first alerting the owner via a text message or some other form of notification.
2. All control and data permissions must be actively given. Therefore all permissions must default to negative, until the resource gives each specific permission. The customer does not relinquish any permissions or privacy just by installing the DER control hardware, and all of the privacy/permissions do not need to be given at once.
3. The DER commands and privacy of the DER usage data must be unquestionably secure.
4. The DER owner must be able to exit the program at any time, and/or switch to another vendor with minimal effort.

The DEGC platform provides the platform to allow privately-owned, heterogenous DER to participate in automated demand response/virtual power plant programs, without compromising security or privacy. The DEGC platform is designed from the ground up to provide these interoperability, security and privacy features.

Beyond customer adoption, the security features are critical for the cybersecurity of the grid. If demand response and virtual power plants experience widespread adoption in the coming decades, the participating DER will become a target for cyberattacks. Thus, the security of DER participating in demand response and virtual power plant programs will be critical for the safe operation of the grid.

DEGC Business Model

People and businesses that are interested in saving money on their electricity bill can purchase DEGC hubs, smartplugs and/or DEGC-compliant appliances/DER. The customer/DER owner is looped into the control decisions by the DEGC phone app. The customer can opt in to the DEGC DER control program on the phone app, and provide the control permissions that they are comfortable with. The DEGC DER control program generates revenue by participating in demand response markets (and/or other markets, as applicable). The market participation

includes customer engagement via the phone application, but does not require the customer to manually adjust any of their DER. A majority share of the revenue generated from the market participation is passed to the DER owners, and a fraction retained by the DEGC system operator to support the DEGC infrastructure, and as profit. The DEGC security structure allows the DER owners to maintain privacy and exit the program at any time.

In an alternative business model, the DEGC system does not interact with any markets directly. Instead, it connects the privately-owned DER with a third party market participation vendor that coordinates the demand response or virtual power plant actions. The DEGC platform allows the DER owners to easily switch between third party market participation vendors, without having to worry about device security or data privacy. With this business model, the DEGC platform takes a smaller share of the revenue that is generated by the market participation.

The Berkeley team has identified a Small Business Innovation Research (SBIR) Phase I grant as a next step for commercializing the DEGC platform. We believe the DEGC platform, and the accompanying market opportunity that we have described, is a strong candidate for the SBIR program since DEGC is based on fundamentally new technology that has the potential to provide significant value to customers.

Hardware in the Loop (HIL) Testing and Simulation

Subtask 4.1: Draft Protocols (M1-M12): Develop Testing Protocols and Performance Criteria for HIL Simulations over varying load (25-150%) and PV penetration (0-150%) levels (relative to peak load). These protocols will specify how a diverse range of test cases is to be created and evaluated in the HIL simulation environment.

Subtask 8.1: Draft Protocols (M13-M24): Develop additional Protocols and Performance Criteria for HIL simulations to produce realistic and challenging scenarios for PBC at increasing scale. These protocols will specify how a diverse range of test cases is to be created and evaluated in the HIL simulation environment. Testing protocols will aim to combine physical devices and data from physical circuits at LBNL with interconnected virtual circuits through OPAL-RT.

The team developed an extensive HIL Testing Playbook that detailed the sequence of test cases to be created and exercised. Each test was characterized by the following:

- a specific distribution circuit model to be load and simulated in OPAL-RT,
- populated with specific generation and loads, including a script for disturbances;
- a control objective with one or more specific performance nodes on the circuit, and
- placement of one or multiple actuators on the circuit, which could be either inverters or load banks.

The team successfully followed the prioritized Playbook list until its available time to experiment in the FLEXGRID was exhausted. To adapt to Covid-19 related constraints in the summer of 2020, which limited the personnel that could work on-site, the team placed greater emphasis on Controller-in-the-Loop (CIL) testing that could be performed in an entirely remote manner. Some protocols were quickly re-written accordingly. The team was very pleased to be

able to complete HIL testing with physical devices for a range of scenarios on four distinct feeders.

Subtask 8.2: Device requirements (M13-M24): Identify devices for HIL testing, including 3 smart inverters connected to the Ametek MX30 Programmable Power Source and various controllable loads connected to the 120V LBNL grid. Measurements from the physical devices under HIL test will combined to map to 100 nodes in the HIL simulation. Plan to install, test, and commission HIL devices not already on location.

Devices

The project team was able to use an existing set of inverters, batteries, and controllable load banks at the FLEXGRID. Fig. 19 shows, from left to right, the rack-mounted OPAL-RT simulator below a screen displaying Power Standards Lab (PSL) micro-PMU measurements, a three-phase Eaton transformer, various disconnect switches, and three SolarEdge inverters. Fig. 20 shows Maxime Baudette, who saw through the project's HIL testing during the Covid-19 crisis as the sole team member physically on site, alongside three micro-PMUs and Tesla Powerwall batteries.

Figure 21: FLEXGRID HIL testing setup. © 2010-2019 The Regents of the University of California, Lawrence Berkeley National Laboratory. Photo Credit: Thor Swift.

Figure 22: Maxime Baudette with micro-PMUs at the FLEXGRID. © 2010-2019 The Regents of the University of California, Lawrence Berkeley National Laboratory. Photo Credit: Thor Swift.

Figure 23: Load racks at the FLEXGRID.

Controllable loads were assembled in seven racks (4ft x 4ft), shown in Fig. 23. These included 30 resistive on/off heaters ranging from 10-150W, and 60 blower motors controlled continuously 0-100% with variable frequency drives (VFDs), ranging from 80-165W. Though the individual loads numbered close to 100, they were aggregated somewhat like loads in a commercial

building with a building management system, which is likely more representative of demand-side resources participating in PBC than small, isolated residential loads.

The original goal of mapping controllable devices to 100 distinct circuit nodes in simulation proved to be too ambitious within the limited project budget. In the team's estimation, the expense of setting up so many individual devices with distinct inputs to the OPAL-RT simulator would have consumed enormous effort while offering little additional scientific value.

Rather, we understood the intent of the task as demonstrating the general scalability of the control framework, both in terms of theoretical controllability and practical implementation. The former issue, controllability, was addressed by *in silico* simulations with many actuators. It should also be noted that one of the circuits used in HIL testing included simulated variable PV generation and loads at hundreds of nodes. The latter issue, implementation, was addressed in the design of the DEGC platform that facilitated all the communications among PBC assets. Although we did not set up a physical experiment with 100 devices independently sending and receiving data, the platform is specifically designed to allow easy scaling to many thousands of participants or devices, abstracted as "processes." Based on the smooth performance of PBC with roughly a dozen devices on the minimalist physical infrastructure built, the team is confident that viability at scale, with increasing numbers of devices and nodes, is not a substantial risk.

Modifications

FLEXGRID was originally built with three separate single-phase PV inverters, connected in a delta configuration on the primary side of interconnection with the LBNL grid or the Ametek grid simulator. For HIL testing scenarios, this configuration would limit the system to be virtually connected to only a single node, as resulting currents of all systems are lumped on the primary side. In order to support more flexible testing scenarios, the transformer's primary side was re-wired to a wye connection, which allows the inverters to be considered independently – for example, as three unbalanced phases or three separate nodes within the model.

Additional modifications included the following:

- A switch logic was implemented for virtually relocating HIL devices on the feeder model, to reduce setup time and allow running more scenarios.
- The Ametek sub-model was updated to reflect phase angle variations, rather than enforcing 120^{o} separation among phases.
- The OPAL-RT chassis was expanded with additional Analog and Digital Inputs and Outputs.
- The Inverter API was modified to reduce latency from 15-30 sec. to less than 3 seconds.
- Three physical micro-PMUs were added to the HIL test setup. These could have been emulated virtually, but using physical μPMUs helped to ensure C37 signal and time stamping compatibility for phase angle differences.

Of all the issues addressed in creating the functional HIL testbed, the inverter API is likely the most problematic. Limitations of standard inverter APIs, in terms of their ability to receive and quickly execute power commands, are a substantial barrier to field implementation and rapid scale-up of PBC.

Subtask 8.3: Test plans (M22-M24): Develop communication and control system architecture and test plans for HIL devices at FLEXLAB. Specifically, for controllable loads located at LBNL, document approach for integration with OPAL-RT ePHASORsim. For smart inverters at FLEXLAB document approach to control with DG-IC controller, and integration requirement for communication with OPAL-RT ePHASORsim.

The communication and control architecture for HIL testing, including the various assets and communication protocols, is illustrated in Fig. 24.

L-PBC communicates (i) with the S-PBC to provide local status and telemetry, and receive reference phasors and local phasor targets, and (ii) with controllable DERs – or, more specifically, to the devices managing those DERs (such as the inverter on a solar panel, the battery controller, or a load controller). In a commercial installation, L-PBC would reside in the field, co-located with DERs. At the very least, the L-PBC needs to receive telemetry and flexibility information to know the current operating state of each device and how much it can be modulated. To control the device, the L-PBC must provide real and reactive power (P and Q) commands, or other instructions (on/off, operating mode, etc). Finally, the L-PBC needs to receive voltage phasor feedback data from local PMUs at sub-second speeds to make real-time device modulation decisions.

The project used physical micro-PMUs to measure voltage phasors at electrical nodes or equipment locations. The µPMUs provide the phasors to the L-PBC in real-time. A copy of µPMU measurements can be archived separately (e.g., sent to a cloud-based instance of BTrDB or PredictiveGrid) or, down-sampled if appropriate, through the L-PBC and S-PBC controller, on a near real-time basis.

The DEGC platform simplifies the configuration of how the above components communicate: the developer simply names the resources required by each controller or process, and the DEGC platform routes the required information to the requester after verifying that the requester is permitted to receive that information. For example, the L-PBC implementation can be configured to receive µPMU data, S-PBC targets and other information from any source, independent of where that source is located on the network. This simplifies testing and means that only the configuration needs to change when deploying the code in a production context.

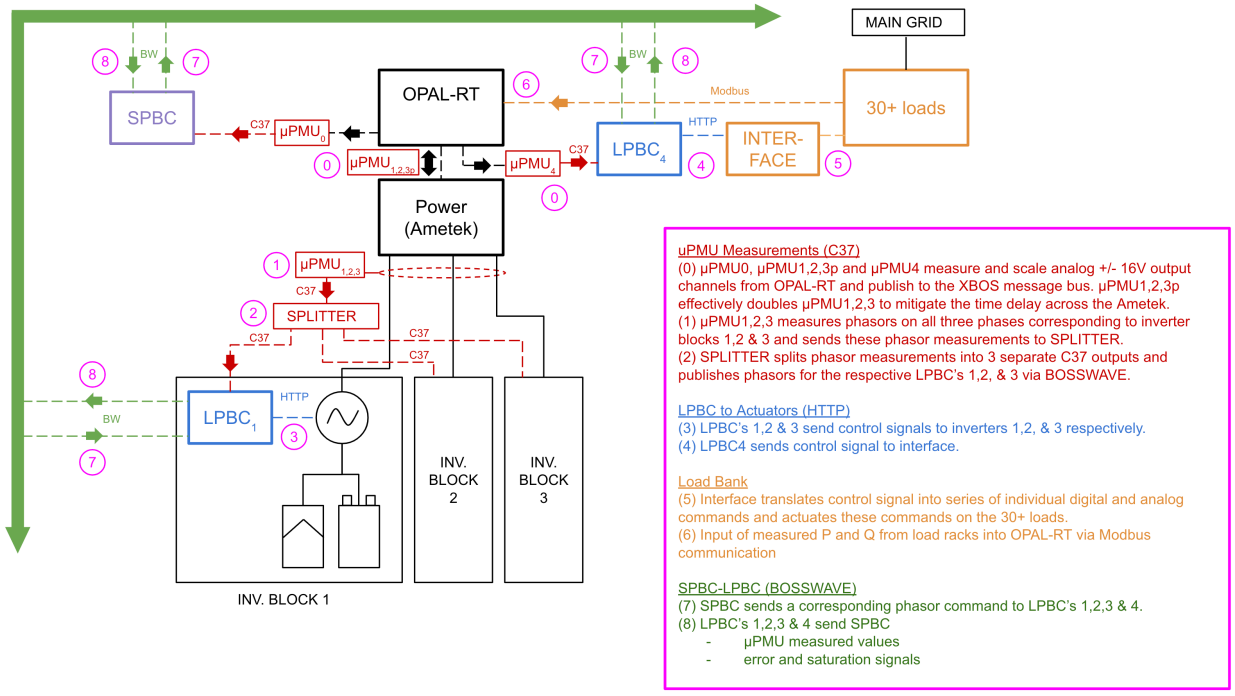

Figure 24: HIL Architecture Schematic.

Subtask 8.4: Cybersecurity and Interoperability Plans (M9-M12): Update the cybersecurity and interoperability plans for PBC platform to include identification of most relevant components for verification.

The Cybersecurity and Interoperability Plans were updated to reflect characteristics of the DEGC platform. Fig. 25 from the original version remains relevant, as it highlights the various layers using the nomenclature from the FOA.

Subtask 8.5: Market Transformation Plan (M9-M12): Update the market transformation plan to include commercialization ideas with known or perceived barriers to market penetration, and product distribution.

The Market Transformation Plan was updated.

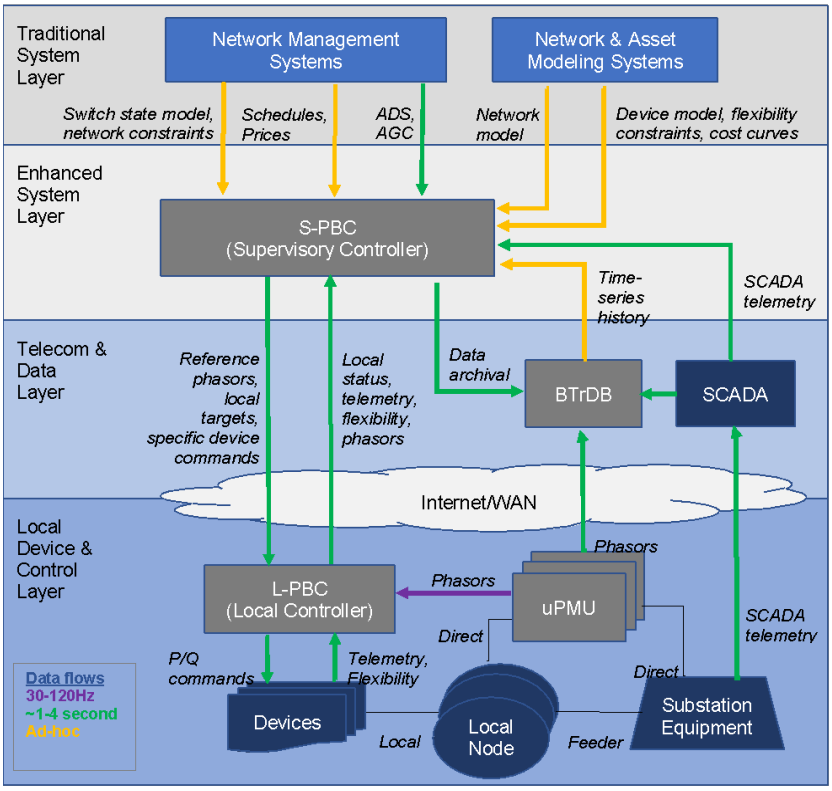

Figure 25: Layers under consideration for cybersecurity and interoperability.

Subtask 11.2: Device requirements (M25-M27): Finalize test and commissioning plans for HIL devices at a physical facility. Given testing and simulation is expensive in physical facilities, these plans will undergo a final iteration with each partner before devices are physically installed.

This subtask was completed early.

Subtask 11.3: Test plans (M25-M27): Evaluate the response of each HIL device to communication and control signals at a physical location, as outlined in the commissioning plans developed in subtask 4.2.

This subtask was completed as part of the test preparations.

Subtask 11.4: Testing (M25-M27): Perform HIL testing on individual and then multiple devices at a physical location.

The team completed HIL testing in two rounds, March and June 2020, despite LBNL site access restrictions due to Covid-19. With thoughtful support from Berkeley Lab management, a modified setup allowed us to accomplish this with one person physically present at the FLEXLAB, and three others coordinating remotely.

Subtask 11.5: Simulations (M25-M27): Evaluate the performance of PBC at a physical location starting with a simulation environment and scaling up to include HIL devices (e.g., smart inverters, batteries, and controllable loads).

Simulations spanned a variety of control scenarios on four different circuits, including a realistic PG&E distribution feeder: IEEE13 balanced, IEEE13 unbalanced, UCB33, and PGE PL0001.

The team was very pleased with the results, which closely matched controller-in-the-loop simulations (minus artifacts that were successfully debugged).

Specific PBC challenges tested included the following:

- Inverters recruited to reject large disturbances from time-varying loads on one, two or three phases.
- Assigned target undergoes a large step change.
- Load racks recruited to track phasor at a different node, away from actuators.
- Load racks used to create disturbances on the feeder, inverters recruited to mitigate.
- Multiple actuators at different nodes recruited to track same target.
- Unrealistic phasor target causes actuators to saturate, propting a revised target request by L-PBC to S-PBC.
- Feeder with low X/R ratio departs from the expected relationship between real and reactive power vis-à-vis voltage magnitude and angle. This challenges the SISO PI controller, but not the MIMO LQR controller.

The simulations included varying percentages (25% to 150%) of PV penetration through the various scenarios on different feeders. As expected, high PV penetration is not a particular challenge factor for PBC performance, except insofar as it creates the potential for large step changes. PBC performance under high generation or net load volatility was meaningfully tested in large disturbance scenarios (e.g., an instantaneous change of 50% of total feeder load) and with large minute-wise variance in the feeder load profile.

The results proved the feasibility of PBC with PI and LQR controllers on four different test feeders with HIL. This greatly expanded on the results from the first phase, which only showed one and two-phase actuation on the IEEE 13-node balanced feeder. We tested the controller's ICDI ("I Can't Do It") capability, multiple actuators, and PBC on a 'physical' feeder (PGE PL0001).

Between the first and second round of HIL testing, various error sources (scaling, coupling, and grounding) were identified and corrected, leading to improved performance. Most importantly, the delta-to-wye modification of the inverter connectivity allowed three-phase control across the different feeders.

LQR results generally demonstrated much higher performance disturbance rejection in terms of convergence time (typically one control iteration). This rapid disturbance rejection matched the performance of the LQR controller during *in silico* simulations. The LQR controller can reject disturbances rapidly because the LQR controller contains an internal disturbance-rejection loop, which is built using an internal network impedance model.

We also tested a live version of the iterative S-PBC algorithm. There were no surprises, which is good news. This test demonstrated that using our linearization techniques, the supervisory controller can run power flow optimizations for a single distribution circuit on a time scale commensurate with the requirements of online operations in near real-time (i.e., on the order of seconds, while the supervisory control time step is intended to be on the order of minutes).

We tested two different local controllers, PI and LQR. It was impractical to deploy the RCAC controller code on location at Berkeley, owing simply to the mechanics of transferring the University of Michigan's code, not because of any inherent incompatibility of the controller design. However, the agreement between *in silico* and HIL simulations for the other two controller types inspired confidence in the team that RCAC would pose no unique challenges in the physical test environment, since the actuation commands this controller sends would be similar to those of the other controllers.

It is difficult to choose highlights from among the many tests performed, to fit within this report at reasonable length. Three examples are discussed here.

The following figures illustrate a series of tests on the IEEE 33-node feeder, with three solar inverters actuating at node 18 to reject a large disturbance at node 26, and maintain a steady voltage phasor difference between the feeder head (node 1) and the performance node (6). The disturbance is caused by physical load banks. This is a challenging scenario due to the location of the inverters, and the low X/R ratio of this feeder. The LQR controller shows smooth disturbance rejection.

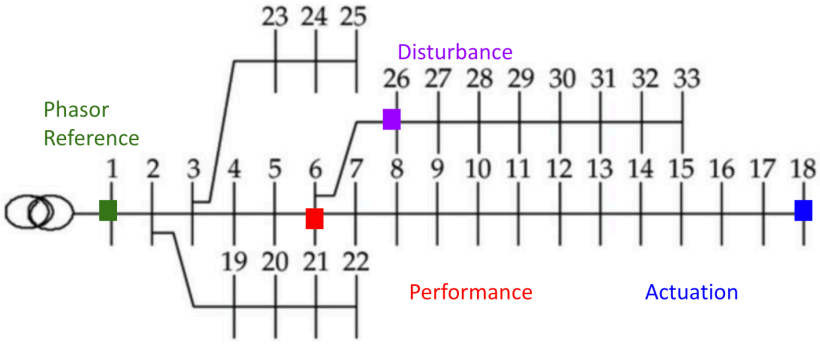

Figure 26: IEEE 33-node feeder for PBC test scenarios (8-1 and 8-2).

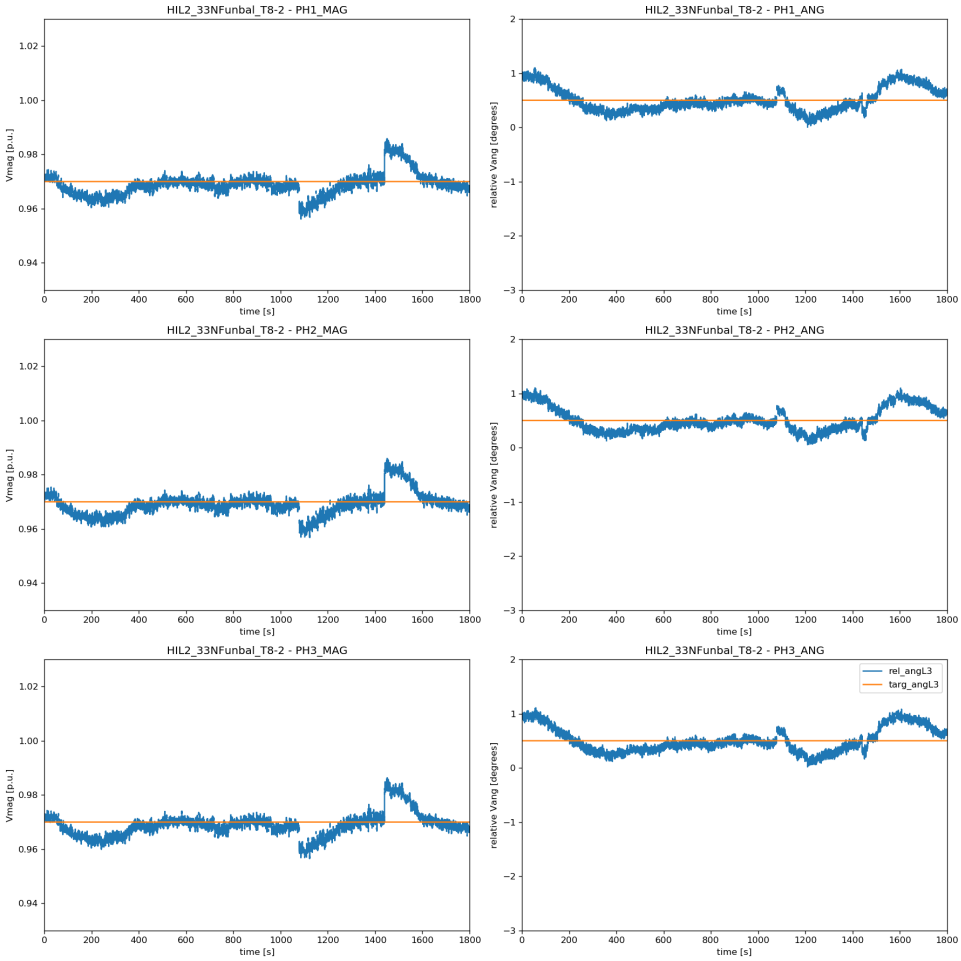

Figure 27: Results for Test Scenario (8-2) with PI controller.

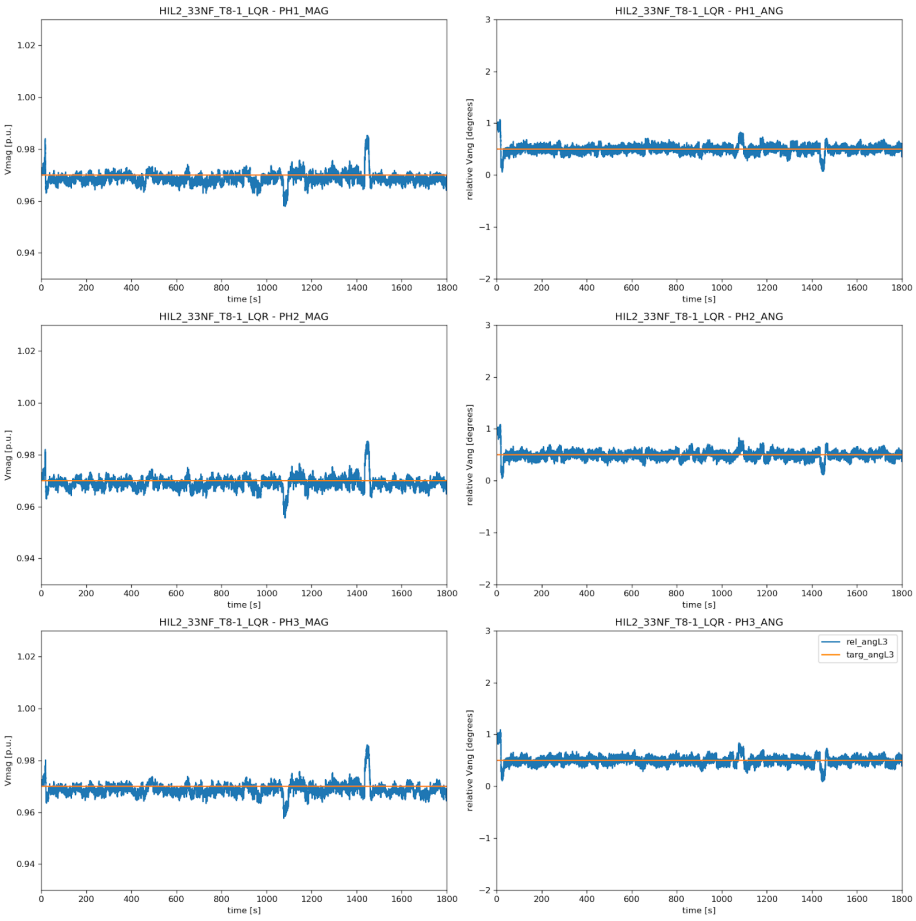

Figure 28: Results for Test Scenario (8-1) with LQR controller.

Test results for the 341-node circuit PL001 are illustrated in Figs. 29 (PI controller) and 30 (LQR). (No diagram is available for this circuit.) This large PG&E feeder features some very unbalanced loads on different phases, with high second-wise variance that makes for an excellent, realistic controller challenge. The test scenario is a single co-located performance-actuation node. Note that the results reflect the load volatility. This feeder is generally harder to control, with big oscillations and a long initial convergence time. However, this near worst-case scenario still yields to PBC without any adverse effects.

The early oscillations in Fig. 30 result from a subtlety in initiating the LQR controller, which has both an error-feedback and a disturbance cancellation. If both are started at the same time (rather than waiting an additional timestep to start the error feedback), then the initial error is double-counted and the controller oscillates initially. This phenomenon was readily eliminated with minor code modifications, but not before the end of the HIL testing period.

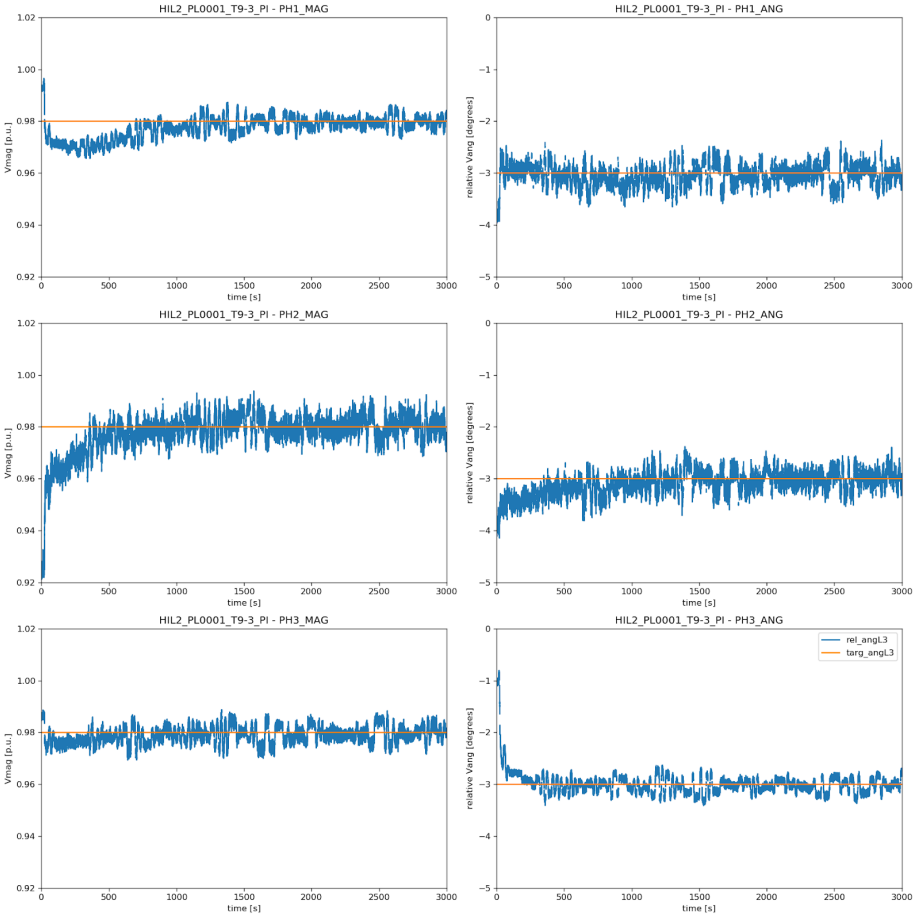

Figure 29: HIL test results (Scenario 9-3) on feeder PL001 with PI controller.

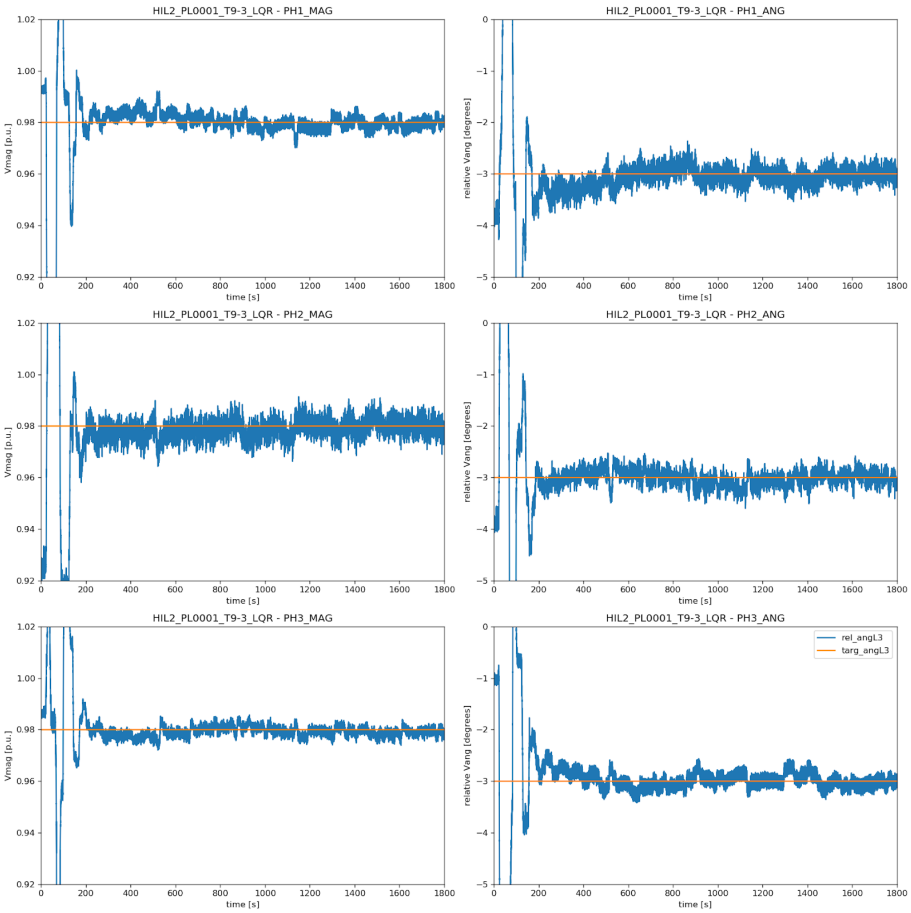

Figure 30: HIL test results (Scenario 9-3) on feeder PL001 with LQR controller. The early oscillations were eliminated in silico, subsequent to the HIL testing period.

Finally, Figs. 32 and 33 illustrate the LQR controller's ability to reject disturbances very quickly, snapping back to the phasor target rather than the decaying exponential behavior that is typical of the PI controller. These tests were performed on the IEEE 13-node feeder, with co-located actuation and performance at node 675, as shown in Fig. 31.

Fig. 33 shows a zoomed-in view of the same test. The control time step is 15 seconds, owing to limitations of the inverter firmware, and the horizontal axis is in units of seconds. Although the inverter response is not consistent with the original performance target of phasor tracking on a one-second time scale, the test results show that there is no shortcoming of the phasor-based controller, which essentially tracks within a single time step.

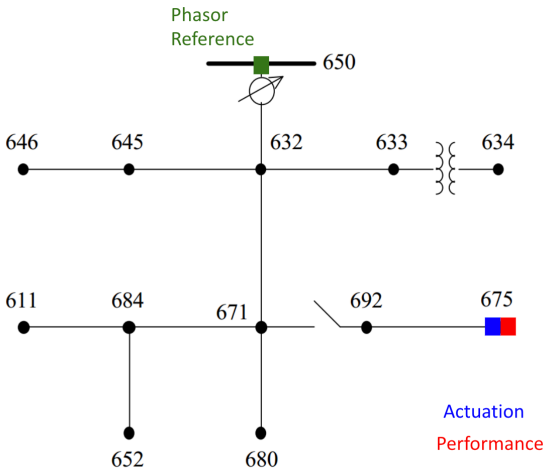

Figure 31: Test scenario (3-3) on the IEEE 13-node feeder.

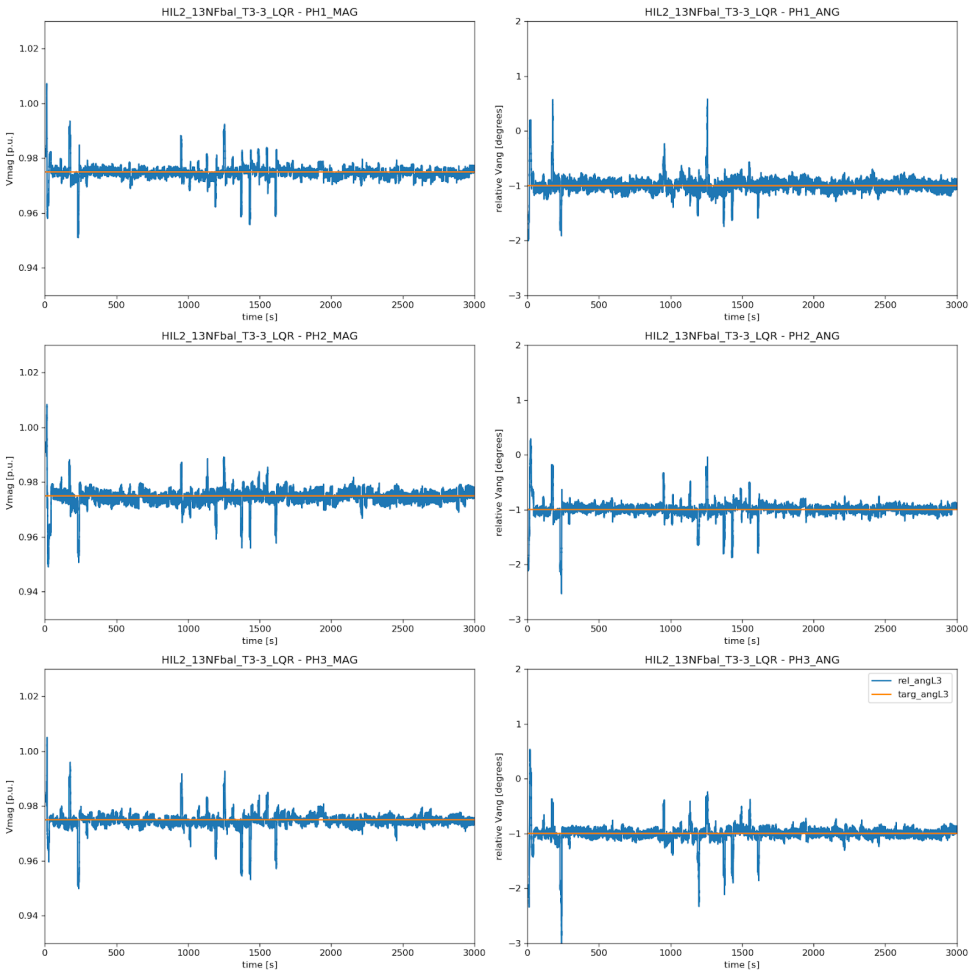

Figure 32: LQR performance in test scenario 3-3.

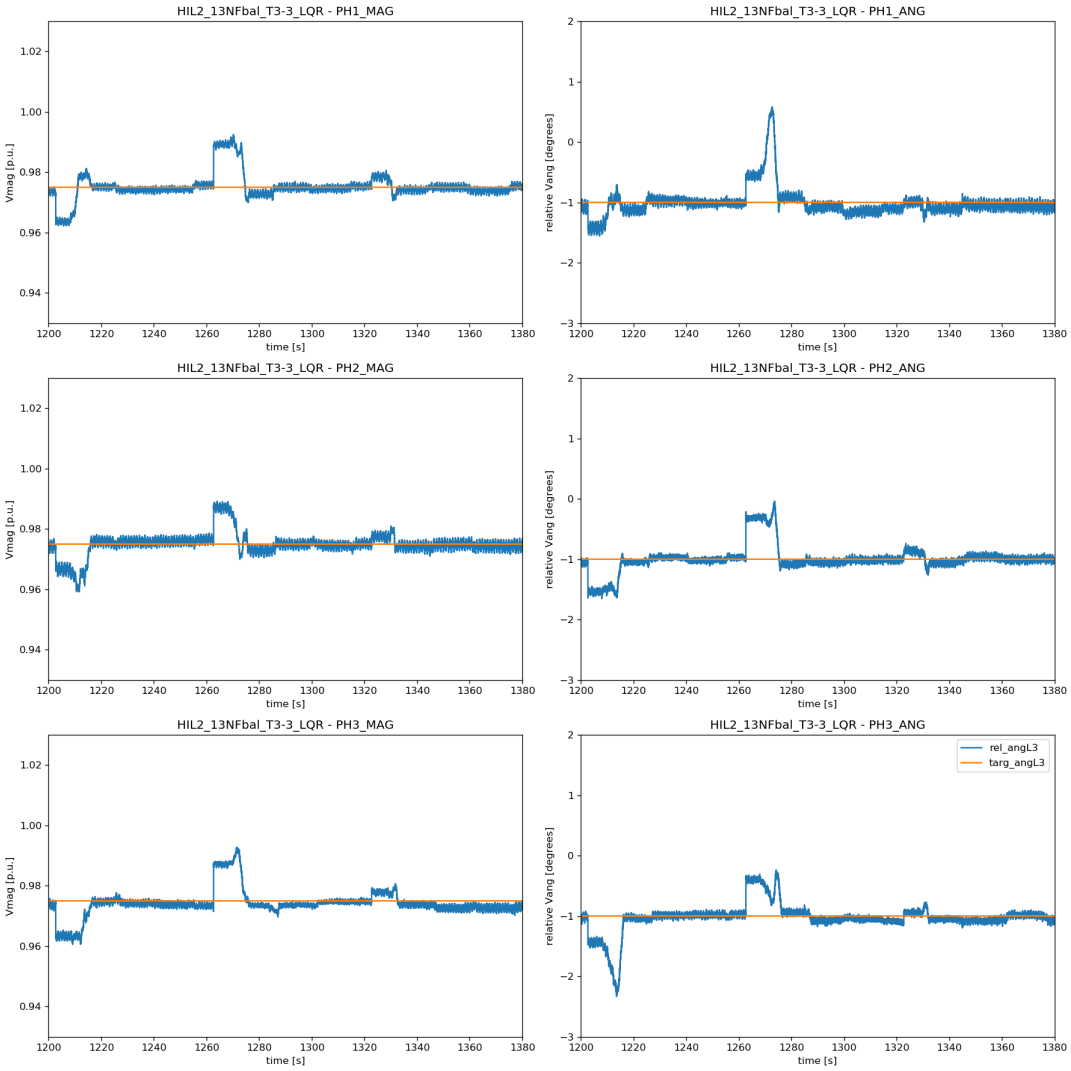

Figure 33: Zoomed-in view of LQR phasor tracking.

In sum, HIL testing confirmed that PBC can serve as a robust control method even in the face of severe, contrived disturbances and challenges. Although the overall speed is limited by the actuation capability in the hardware, these limitations created no instabilities or other adverse effects. With one exception at the beginning of the control interval on the PL0001 circuit, we find PBC controllers driving the system toward the given target.

Subtask 11.6: Final Cybersecurity and Interoperability Plans (M28-M30): The cybersecurity and interoperability plans will be included in the final report, including the results of the verification tests.

The Final Cybersecurity and Interoperability plan is attached to this report as a separate document due to its 21-page length. It was updated from the original version by accounting for the new DEGC platform, based on the XBOS and WAVE data infrastructure, developed by UC

Berkeley. A key property of this platform lies in its security features based on decentralized authentication.

Subtask 11.7: Final Market Transformation Plan (M28-M30): The market transformation plan will be included in the final report with proposed commercialization timeline, financing, product marketing, competitors, distribution channels, and potential vendor identifications.

The Final Market Transformation Plan is attached to this report as a separate document due to its 18-page length. Next steps by the project team, in addition to those discussed under Subtask 10.5, include the following:

Field Demonstration

In the context of the Oakland EcoBlock project funded by the California Energy Commission and led by CIEE, the Berkeley team may have an opportunity to exercise PBC in controlling a block-scale microgrid with ca. 120 kW solar PV, battery and inverter capacity on the basis of phasor measurements on the distribution feeder. Located on a cul-de-sac with a single-phase lateral, this islandable microgrid will use the extant 4-kV infrastructure and interconnect to the 3-phase main feeder on Fruitvale Avenue. Construction is anticipated to begin in late 2021. The project team is actively seeking out funding opportunities to support a comprehensive PBC demonstration to leverage this unique opportunity.

Trade Show Demo

One key next step identified by the Market Transformation Plan is to educate the prospective customer base about the capabilities of PBC, and potential applications to their particular systems. To this end, OPAL-RT created a demonstration to include in their repertoire for trade shows. This involves a real-time model to demonstrate the phasor-based controller developed by the UCB team.

Specifically, a local phasor based controller (L-PBC) was implemented in the ePHASORsim PL0001 distribution feeder model. Ten nodes in the distribution feeder were selected for the L-PBC control. The objective of the demo model was to show that the L-PBC can achieve the phasor set points in the distribution nodes by actuating the inverters in the distribution feeder. In the demo model, it is assumed that the S-PBC controller will send the phasor voltage targets for those selected nodes also called performance nodes, and based on the targets the L-PBC controller will change the active and reactive power of the inverters collocated at the performance nodes. For the demo, the phasor setpoints were fixed in the simulation assuming the targets were sent by the S-PBC controller.

For the demo, a LabVIEW panel was created to interact with the demo model. The RT-LAB model can interact with the LabVIEW panels via an API. The LabVIEW panel in the demo provides two functionalities, one for the user controls and another to visualize the results from the real-time model. In the control section, the user can turn ON/OFF the L-PBC control. With the controller ON, the active and reactive power of the inverters are controlled to meet the phasor targets in the ten nodes. With the controller OFF, the action of the inverters is turned off. Additionally, the user can select different time-intervals: morning, afternoon, and late afternoon to test the L-PBC control with different load and PV profiles. The panel also provides the plots

for the inverter active, reactive power actuation as well as the load and PV profiles during the simulation.

In the demo model, varying PV and load profiles were used resulting in the change in node voltages in the PL0001 distribution feeder. Because of this, without L-PBC activated, the node voltages could not follow the phasor targets set by the supervisory controller, which is defined as constants for the demo purpose. The phasor setpoints used in the simulation are shown below each phasor plot and can be changed in real-time to mimic the target set points being changed by the S-PBC. The ten voltage phasor plots for the performance nodes show the reference and measured voltage phasors at the selected nodes. Each phasor plot is equipped with a LED indicator that turns green or red. If the measured node voltage is within the 2% and 2 degrees of the phasor magnitude and angle setpoints, respectively, the LED turns green, indicating the target setpoints are met; else, the LED turns red. The demo illustrates that a mismatch between the node voltage and phasor targets because of the varying load and PV penetration, but with L-PBC, the voltage phasor can be changed to track the phasor targets as set by the S-PBC controller.

Figure 34 shows a test case with the morning load and PV profile. In the simulation, as can be seen in the figure, there is mismatch between the measured voltage phasors and the target voltage phasor when the L-PBC is off. As the L-PBC controller turns on, the voltage phasors at the nodes gradually change to match the voltage targets as seen in Fig. 35.

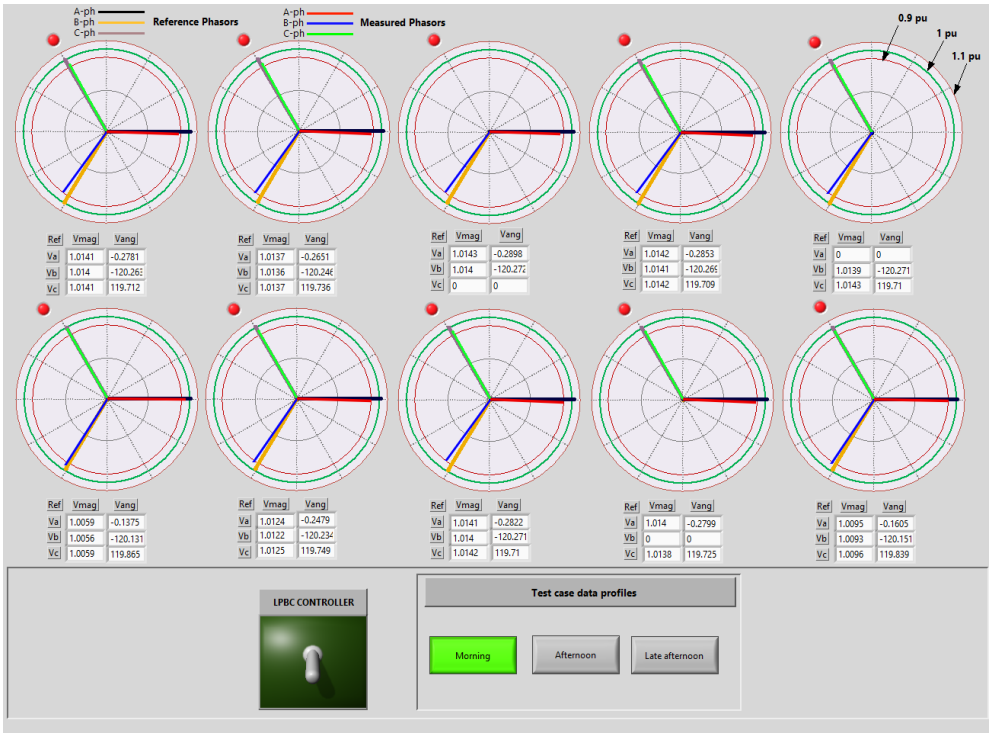

Figure 34. Morning case demo with L-PBC off.

Figure 35. Morning case demo with L-PBC on.

DEGC Platform

Another emerging theme is the potential for broader application of the data infrastructure developed for this project, irrespective of the PBC control paradigm per se. Data collection and management expanded to real-time analytics will play an increasingly important role in the power industry. Our Distributed Extensible Grid Control (DEGC) platform dsicussed under Subtask 10.5 above is adaptable to many uses that rely on some integration of real-time sensor data streams, analytics and control. It is particularly apt for any approach involving multiple parties, as in the actuation of diverse distributed resources, where issues of restricted grid data access and authority come into play.

Subtask 11.8: Value Analysis: A value analysis that shows the net benefits of the project solution, the benefits' sustainability at levels of penetration of solar greater than 100% of peak load. This analysis will be included in the final report.

The project completed a value analysis for the use case of reliability support at the transmission level. The full-length Value Analysis Report (34 pages) was submitted in August, 2020.

With input from PG&E and Dominion Energy, the team previously determined that the most economically compelling use cases for PBC would be motivated by grid security considerations, rather than cost-benefit analyses at the scale of individual distribution feeders. Although it is strategically important that PBC supports very high penetration levels of variable solar PV generation, we take this as a prerequisite rather than an explicit selling point, since it is difficult

to assign an economic value to the utility from increasing feeder hosting capacity. Similarly, the ability to improve power quality on distribution circuits suffering from high volatility in either generation or load is hard to monetize directly in the distribution context. However, extrapolation of the impact of PBC-controlled distribution circuits on the transmission level captures tangible economic benefits.

For this analysis, we created specific scenarios where planned transformer or line outages create conflicts with N-k security constraints, forcing sub-optimal operation. An algorithm performs an exhaustive search for N-k insecure states on a 14-bus transmission test system. Multiple contingency scenarios were analyzed. We assume that high penetration levels (>100%) of solar PV generation are present behind each substation, along with storage and load control, and that an individual distribution feeder can be recruited through PBC to modulate its net real and reactive power demand relative to the transmission system.

In each case, it was possible to remediate voltage and power flow violations with a PBC performance node in the vicinity. Using value of lost load (VOLL) and timing on equity analyses for scenarios with a single planned system upgrade, we found that mitigation of security constraints by PBC in the scenarios studied can yield a present value in the range of $1.4 to 1.9 million, accomplished by a single, well-placed PBC performance node at a substation.

This is an encouraging result, since the value far exceeds the cost of hardware deployment (presently, on the order of $10k). Although the scenarios were specifically picked for highlighting potential value and do not represent an average, the deployment of PBC infrastructure would generate additional value by enabling multiple parallel applications, including the distribution-specific test cases studied in this project.

Products Developed and Technology Transfer Activities

Publications

1. A. Ul Islam, E. Ratnam and D. Bernstein, “Phasor-Based Adaptive Control of a Test-Feeder Distribution Network.” IEEE Transactions on Control Systems, 2019.
2. A. von Meier, E. Ratnam, K. Brady, K. Moffat and J. Swartz, “Phasor-Based Control for Scalable Integration of Variable Energy Resources.” *Energies* 2020, 13(1), 190. https://doi.org/10.3390/en13010190
3. K. Moffat, M. Bariya and A. von Meier, “Real Time Effective Impedance Estimation for Power System State Estimation.” IEEE Innovative Smart Grid Technologies (ISGT) Conference, Washington, DC, Feb 2020.
4. J. Swartz, T.G. Roberts, A. von Meier and E. Ratnam, “Local Phasor-Based Control of DER Inverters for Voltage Regulation on Distribution Feeders.” *IEEE GreenTech Conference*, Oklahoma City, OK, April 2020.

5. K. Moffat, M. Bariya and A. von Meier, “Unsupervised Impedance and Topology Estimation of Distribution Networks—Limitations and Tools." *IEEE Transactions on Smart Grid* 2020, 11(1).
6. G. Fierro, K. Moffat, J. Pakshong and A. von Meier, “An Extensible Software and Communication Platform for Distributed Energy Resource Management.” IEEE SmartGridComm'20, November 11-13 2020.
7. K. Brady and A. von Meier, “Iterative Linearization for Phasor-Defined Optimal Power Dispatch.” North American Power Symposium (NAPS), Tempe AZ, April 2021 (accepted).
8. J. Swartz, B. Wais, E. Ratnam and A von Meier, “Visual Tool for Assessing Stability of DER Configurations on Three-Phase Radial Networks.” Submitted to IEEE Powertech 2021. arXiv preprint available at arXiv:2011.07232
9. K. Moffat, J. Pakshong, L. Chu, G. Fierro, J. Swartz, M. Baudette, C. Gehbauer and A. von Meier, “Phasor-Based Control with the Distributed, Extensible Grid Control Platform.”
10. M. Baudette, L. Chu, C. Gehbauer, K. Moffat, J. Pakshong, J. Swartz and A. von Meier, “Hardware in the Loop Benchmarking for Phasor-Based Control Validation.” (in preparation)
11. K. Moffat and A. von Meier, “Local Power-Voltage Sensitivity and Thévenin Impedance Estimation from Phasor Measurements.” (in preparation)

Open-Source Code

PBC Feasibility Tool: https://github.com/jaimiosyncrasy/heatmap

DEGC Repositories: https://github.com/gtfierro/DEGC

- Message bus: https://github.com/immesys/wavemq
- Auth platform: https://github.com/immesys/wave
- Main DEGC implementation: https://github.com/gtfierro/xboswave/
- DEGC Software Framework: https://github.com/gtfierro/xboswave/tree/master/python/pyxbos
- Quickstart for PBC: https://github.com/gtfierro/energise-quickstart