

Distribution Feeder-Scale Fast Frequency Response via Optimal Coordination of Net-load Resources—Part II: Large-Scale Demonstration

Blake Lundstrom, *Senior Member, IEEE*, Sourav Patel, *Student Member, IEEE*, and Murti V. Salapaka, *Fellow, IEEE*

Abstract—This paper is the second of a two-part series in which we develop and experimentally demonstrate a hierarchical control solution for optimally coordinating thousands of deferrable loads and distributed energy resources (DERs) to provide fast frequency response (FFR) from an entire distribution feeder. In Part I, we developed and proved practical algorithms for fast, cost-based optimal dispatch and for determining the optimal amount of headroom to operate solar inverters with to support FFR dispatch while minimizing opportunity cost. Simulation results in Part I demonstrated the advantages of the hierarchical dispatch approach in being able to maintain fast solution times needed for FFR even when the problem size increases. In Part II, we implement the algorithms developed in Part I in a novel, large-scale power hardware-in-the-loop experiment including embedded controllers and more than 100 powered appliance loads and DER connected to a simulated real-world distribution system with more than 10,000 controlled devices. Experimental results from multiple scenarios confirm that the optimal FFR dispatch approach scales well and can optimally coordinate more than 10,000 net-load resources across a distribution network while achieving hardware response times within 500 ms, which is not possible using state-of-the-art optimal coordination approaches.

Index Terms—fast frequency response (FFR), frequency control, optimal dispatch, flexible load, distributed energy resources (DERs), demand response, coordinated response

I. INTRODUCTION

In Part I of this two-part paper [1], we presented background and motivation for providing FFR from optimally coordinated net-load resources (i.e., DER and flexible loads) across a distribution network. A thorough review of related work in this area, including existing approaches for the optimal coordination and provision of primary and secondary frequency response using distributed resources was also presented in Part I. In addition, we designed a solution that can optimally coordinate thousands of heterogeneous DERs and flexible loads connected in an electric distribution network as a virtual power plant to provide FFR to the bulk power system that is dispatchable, predictable, and meets general FFR service performance requirements

B. Lundstrom (e-mail: blake.lundstrom@nrel.gov) is with the Power Systems Engineering Center, National Renewable Energy Laboratory, Golden, CO, USA and the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN, USA.

S. Patel and M.V. Salapaka are with the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN, USA.

This work was authored in part by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding provided by the Advanced Research Projects Agency-Energy (ARPA-E) under grant no. DE-AR0000701. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

(i.e., ability to achieve ≥ 1 -MW capacity, response within 500 ms, and meet aggregate set point within 1%). Algorithms for fast, cost-based optimal dispatch of net-load resources within a single net-load unit (NLU) were developed, and their optimality was analytically established. We also developed an approach for determining, in real time, the optimal amount of headroom to maintain for a solar inverter within an NLU that is participating in FFR dispatch. The optimal dispatch algorithms were scaled up to achieve near-optimal dispatch of net-load groups (multiple NLUs or buildings connected to the same distribution transformer) and an entire distribution feeder with 800 net-load groups (NLGs) representing 40,000 net-load resources. Simulation results from thousands of simulation runs were presented to validate that the developed dispatch approach achieves results within 1% of optimal (within 0.3% for the majority of participation levels) while maintaining consistently fast solution times applicable to FFR, regardless of problem size.

In this article, the solution for FFR from optimally-coordinated net-load resources developed in Part I is implemented in a large-scale power hardware-in-the-loop (PHIL) experiment with more than 100 powered DER and controllable loads, more than 10,000 simulated net-load resources, and a real-world distribution network model to experimentally demonstrate the scalability and performance of the developed approach.

PHIL simulation is an important tool for derisking modern grid technologies in a controlled laboratory setting in which actual at-power devices and systems are interconnected in software simulations of complex, large-scale power system scenarios. Using PHIL, the closed-loop interaction between the hardware under test (HUT) and the simulated power system can be evaluated under a variety of grid configurations and repeatable operational scenarios without the expense or risk of a field deployment. PHIL techniques have been frequently used to study the interaction of modern grid technologies in power distribution systems. The most common application is to connect a single power device to a simplified model of a power system. For example, a 15 kW motor connected to a simple low-inertia power system was tested in [2], a single photovoltaic (PV) inverter's anti-islanding capabilities are evaluated using a simple power system model in [3], [4], a PV inverter's interactions with a simple power system model with on-load tap changing transformer is considered in [5], a wind energy conversion system connected to a 5-bus power system is tested in [6], and simple P/Q loads on an at-power 25 kV distribution line were connected to a simulated transmission system in [7]. More recently, PHIL

simulation has been used to interconnect multiple HUT at multiple independent points of common coupling (PCCs) in a distribution system model. References [8], [9] evaluated the P/Q control interactions [8] and anti-islanding detection capabilities [9] of up to three independently-connected PV inverters. Two PV inverters' P/Q control capabilities on a low voltage network were evaluated in [10]. In [11], two PV inverters were interconnected to two independent PCCs on a simulation of the IEEE 8500-node test feeder with physics-based loads.

This PHIL experiment completed in this work is unique versus state-of-the-art PHIL experiments because it includes a large number of distributed at-power devices under test (more than 100 individually controllable, powered devices) connected at nine different interconnection points in the simulated distribution system model. In addition, the experiment included a real-time simulated electric distribution system model with over 1000 nodes and more than 10,000 cosimulated net-load resources. We are unaware of any other PHIL demonstration to date including this many powered devices and individual PHIL interconnection points.

An overview of the PHIL experiment design is presented in Section II, results from a PHIL experiment representing an NLG with nine residential buildings are shown in Section III, and results from the full-scale electric distribution system-scale PHIL experiment with coordinated FFR are given in Section IV.

II. OPTIMAL FFR DISPATCH PHIL EXPERIMENT PLATFORM

A novel PHIL experiment platform (see Fig. 1) incorporating more than 100 powered DER and loads, a real-time simulated electric distribution system model, more than 10,000 cosimulated net-load resources, and a hierarchical control solution deployed on embedded hardware platforms and real communications infrastructure is developed. This allows for the optimal FFR dispatch approach developed in Part I to be scaled and tested in a realistic electric distribution system simulation. Further detail on the components of the platform are presented in the following subsections.

A. Power Hardware Under Test

The over 100 powered HUT are organized into ten hardware groups (HGs), which are described in this section.

1) *Laboratory Homes (HG1–HG2)*: Each lab home contains four residential appliance loads that are controlled, additional appliance loads that are uncontrolled, and one custom ("Type 4 inverter") 2-kVA inverter system (see Fig. 2). The four controlled appliance loads are connected through load sensing and control interfaces (LSCIs), which provide high-speed measurement feedback and on/off load controllability to the endDevice controller, which is a controller with a common communications interface and generalized structure with specific logic for the various devices controlled (i.e., load, multiload, Photovoltaic inverter, battery inverter, and PV+batt. inverter). In HG1, a refrigerator (Whirlpool WRL767SIAM00), clothes washer (Whirlpool WFL98HEBU0), window air conditioner (Frigidaire FFRE0633S1), and oven (Bosch HBLP651UC) are controlled; and in HG2, the controlled loads

are a refrigerator (GE Profile PSQS6YGY), plug-in heater (Bovado 0CZ449), set of light bulbs (bank of 5 incandescent, 10 compact fluorescent, and 10 LED light bulbs), and oven (Maytag MER8674). A custom 2-kVA inverter system is developed based on a simple H-bridge topology to achieve faster response (within 1–2 AC cycles) to AC power commands than was available via commercially-available residential inverters. The inverter DC inputs are powered by commercial photovoltaic (PV) simulators. More detail on our LSCI and custom inverter implementations can be found in [12]. HG1 and HG2 are connected to the same 240-V point of common coupling (PCC) in the simulated distribution network.

2) *DER Racks (HG3–HG8)*: Each DER rack (e.g., see Fig. 3) contains hardware representing three different residential buildings: (i) a home with deferrable load and a PV string inverter (Type 1); (ii) a home with deferrable load and 12 PV microinverters (collectively referred to as a Type 3 inverter); and (iii) a home with deferrable load, a PV string inverter (Type 2), and a battery energy storage system with inverter (Type 5). All inverters are commercial off-the-shelf (COTS) residential-scale inverters with AC output ratings of 240 V, 3.8–5 kVA and are managed via Modbus TCP from their local endDevice controllers. Battery energy storage systems are also COTS with ratings of 3.3 kW/6.4 kWh (HG3) and 5 kW/9.8 kWh (HG4–HG8). The load in each home is represented by a set of virtual loads (deferrable and nondeferrable) driven by real-world residential load profiles (see Section II-B2 for details). In HG3 and HG4, the output of each home's total virtual load is given as the command to one of six resistive load banks, resulting in physical power consumption for each home's load. In HG5–8, the total virtual load is given as a command to a real-time simulated load connected at the same PCC of the simulated distribution network as the associated physical DER rack. Each DER rack is connected at an independent 240-V PCC (interface points P2–P7 shown in PHIL interface and real-time simulation panels of Fig. 1) in the simulated distribution network model. PV inverter DC inputs are powered by commercial PV simulators.




Fig. 3. Photos of equipment comprising three of the six DER racks (HG3–HG8). Photos by NREL.




Fig. 1. PHIL experiment platform, including control, power hardware under test, PHIL interface, and real-time simulation components (see labels at left)




Fig. 2. Photos of the two lab homes (HG1 and HG2) with residential appliance loads controlled denoted. One Type 4 (custom) inverter (pictured at right inset) is used with each HG. Photos by NREL.




Fig. 4. Photo of equipment comprising two utility-scale PV+battery inverter systems (HG9-HG10). Photos by NREL.

3) Utility-Scale PV and Battery Inverters (HG9-HG10):

HG9 and HG10 each include one utility-scale (100-kW, three-phase) battery inverter and one PV inverter (250-kW, three-phase) that are coordinated to track a commanded AC power output while simultaneously maximizing the use of available input solar energy. This coordinated action is accomplished via an integrating controller (see Fig. 4). Each battery inverter is powered by a 100-kW/33-kWh Lithium-ion battery rack, and the PV inverters are powered by commercial PV simulators. Each HG is interfaced at an independent, three-phase PCC (P8,P9 in Fig. 1) in the simulated distribution network model.

B. Real-Time Simulation

As shown in Fig. 1, the real-time simulation includes an electric distribution network model with additional building and solar device inputs implemented on a digital real-time simulator (DRTS) as well as additional NLG models that are cosimulated in real time on a separate computational platform. For this experiment, a 16-core Opal-RT OP5700 is used as the DRTS and up to (20) 36-core servers are used to cosimulate the NLGs—including their constituent load profiles, devices,

and controllers—in real time. UDP/IP communications are used to interface the two platforms for real-time cosimulation. Here, UDP communication is used for real-time cosimulation as in a real-time simulation, only up-to-date information from the current timestep should be used; if a packet is dropped, we do not want to retry (as in TCP/IP communications) and potentially receive that out-of-date information at a later timestep in the simulation. See [13] for more information on real-time co-simulation using UDP.

1) *Electric Distribution Network Model*: An electric distribution network model based on a real-world system with more than 2,000 customers in Aurora, Colorado is used as the basis for this study. The original distribution network model is augmented with additional single-phase residential and three-phase commercial-scale PV generation to achieve a high-penetration solar case study, as detailed in Table I. Distributed PV systems are added at each residential transformer secondary with individual capacities allocated to follow the statistical distribution of PV system capacities currently in the field while respecting existing transformer capacities. The network model was converted from (and validated against) the native format provided by the electric utility to Opal RT's ePHASORSIM [14] format to allow for real-time phasor domain power system simulation at a time step of 10 ms. The seven residential (P1–P7 in Fig. 1) and two commercial (P8, P9 in Fig. 1) PHIL PCCs are chosen as shown in Fig. 1 to include nodes at a variety of locations. Details on the specific device capacities and mappings are given in Table IV.

TABLE I
DISTRIBUTION NETWORK MODEL CHARACTERISTICS

Location	Aurora, CO, USA
# Customers	~2,000
% Residential / % commercial load	52% / 48%
# Nodes	1,122
Peak load	7 MVA
Annual energy consumption	25,346 MWh
PV inverter capacity	8.06 MVA
Residential	4.96 MVA
Commercial/utility	3.1 MVA
Annual PV energy delivered	12,714 MWh
PV penetration (by annual energy)	50.2%
Battery inverter capacity	1.01 MW
Residential	0.02 MW
Commercial	0.99 MW

The network model's feeder head is connected to a stiff voltage source with source impedance matching that of the substation transformer. Further bulk power system modeling—including the closed-loop impact of feeder power flow on bulk power system frequency—is beyond the scope of this work. The focus of this work is to demonstrate that tens of thousands of distributed net-load resources can be optimally coordinated to provide an aggregate feeder head active power response of desired magnitude and response time commensurate for FFR. The network model provides a realistic power network configuration for interconnection of net-load resources and is used to identify any potential distribution-level impacts of providing FFR (e.g., voltage or power flow beyond desired operating limits). System frequency is determined using an open-loop frequency profile. A frequency profile with an underfrequency step from 60 to 59.9 Hz is

used to more easily quantify FFR dispatched device response times. The framework and methodology are also tested using frequency profiles based on real-world measured data, which further corroborate their effectiveness, though these results are not reported here. All NLU controllers, controlled net-load resources, and power HUT are frequency dependent in this work, whereas the line, load, and transformer models within the distribution network model are frequency independent (a common assumption that doesn't adversely affect the results shown here).

2) *Simulated Devices and NLUs*: The simulated PV inverter, battery inverter, and load devices in the real-time simulation (both in the DRTS and cosimulated models) use constant power injection models fed by solar irradiance and load profile data. All load and uncontrolled inverter device models are frequency independent in this work. Commercial and nondeferrable residential load profiles are determined using load allocation techniques based on transformer capacities and time-series feeder head power data provided by the electric utility. To model realistic deferrable residential loads, 1-Hz power readings from real homes in the Reference Energy Disaggregation Data set (REDD) [15] are used. This data set provides load data at the load circuit level that is time-correlated, driven by realistic user interaction, and at a high enough resolution (1 Hz) to represent fast load transitions as well as longer timescale characteristic profiles. From the REDD, we sample sets of four appliance (e.g., electric heat, dishwasher, water heater, and fridge in one home) loads in five different homes during the same 1-hour period (18:00–19:00) from each day in spring to obtain more than 1,000 sets of (4) time-correlated plug loads. The specific load profiles for the four deferrable loads within a given NLU are selected at random from the 1,000 sets of sampled plug load data, and then a uniformly distributed time offset [-30 min, 30 min] is applied to each profile to maintain appropriate load diversity. Solar irradiance data from the same geographic location as the real-world distribution network and at the same time of day and season as the load data are used.

C. PHIL Interface

The PHIL interface connects the power HUT in HG1–HG10 (see Fig. 1) in a closed loop to the real-time simulated distribution network at $N_j = 9$ different PCCs (P1–P9). For P1, two phases of a three-phase, 45-kVA power amplifier (Ametek MX45) are used to obtain a split-phase 240-V output. For P2–P7, two three-phase power amplifiers are used (one MX45 and one 90-kVA Ametek RS90); each phase of these three-phase power amplifiers is controlled independently and interconnected through a transformer to derive an independent, split-phase 240-V output. Finally, P8–P9 are three-phase PCCs driven by separate 270-kVA three-phase power amplifiers (Ametek RS270) with independent phase control. All five power amplifiers have a -3dB bandwidth of approximately 1,950 Hz. Thirteen sets of voltage (LEM DVL series, 15-kHz bandwidth) and current (LEM LF series, 100-kHz bandwidth) transducers at each PCC phase output (i.e., after power amplifier and, if applicable, the transformer interface) are used to measure instantaneous voltage and current for feedback.

In the PHIL interconnection forward path, root mean square (RMS) voltage, V_j^* , measured at each PCC in the simulated network model and the modeled system frequency, f^* , are used as references to the power amplifiers. A simple proportional-integral control-based voltage compensator drives the measured voltage, V_j , at each physical PCC's output terminals to follow V_j^* , ensuring that the HUT experiences the same voltage as in the network model at that point. The final voltage reference, $V_j'^*$, and frequency reference, f^* , are communicated digitally to each power amplifier. In the PHIL interconnection return path, instantaneous voltage and current measurements at each physical PCC's output terminals are connected to the DRTS's analog-to-digital converter, scaled, and then used to calculate the corresponding active and reactive power for injection into the distribution network model. The PHIL interfacing model is executed at a time step of 200 μ s with communication to the power amplifiers and distribution network model every 10 ms (the time step of the phasor-domain network simulation).

D. Hierarchical Control Platform

As shown in Fig. 1, four levels of control are employed in the hierarchical control platform. The software for this control platform is developed in Python and designed to be deployable in a variety of embedded platforms, from low-cost single-board computers (SBCs) to high-performance servers. At the device layer, an endDevice controller interfaces one or more net-load resources within an NLU to the NLU controller (NLUC). The endDevice controller relays device RMS power measurements and operating limits to the NLUC and translates NLUC dispatch signals into device commands. In the experiment configuration of Fig. 1, the four deferrable loads within HG1 have individual endDevice controllers implemented on Raspberry Pi 3B+ SBCs [16] that also function as the LSCI controller and provide measurements to the NLUC every 50 ms. In HG2, a multiload endDevice controller on a single Raspberry Pi 3B+ manages the four deferrable loads through a circuit-level LSCI. For HG3–HG10, the 50 endDevice controllers, as well as the associated 25 NLUCs and 3 NLG controllers (NLGCs), are deployed on a single 12-core Linux server. The inverter endDevice controllers interface with their physical inverters via Modbus TCP.

At the second layer of control, NLUCs gather measurements from each connected endDevice and maintain a table of assigned device IDs, measurements, capacities, and user-assigned costs that is updated each time a new measurement is received. The NLU periodically (every 100 ms) publishes this table in a standard format to the NLGC, and also translates NLGC dispatch commands and routes them to the appropriate endDevice. Each NLUC and NLGC hosts a local publish/subscribe communication bus (implemented using Zero MQ [17] over TCP/IP) to communicate with their endDevices and child NLUs, respectively. This communication architecture—in which controllers can communicate only one layer up or down—preserves privacy (e.g., NLGs cannot directly address endDevices and do not have knowledge of specific devices) and enables specific communications protocols or security measures to be implemented at each layer as needed.

Building owners interact with the NLUC by setting unit costs and enabling/disabling dispatch for each NLU device.

At the third layer of control, NLGCs gather device status tables from each connected NLU and participate in optimal dispatch calculation. If operating in single NLG dispatch mode, NLGCs calculate an optimal dispatch for only the devices within their NLG using Mthd. 1–3 described in Part I. If operating in global dispatch mode, NLGs calculate the best net-load response options for their NLG using Mthd. 1 and Mthd. 4 (described in Part I), and this information is sent to the Distribution System Master Controller (DSMC) to support calculation of a globally-optimal dispatch solution. After optimal dispatch decisions have been made, the NLGC sends these commands to each NLUC for execution within each NLU. The DSMC and all NLGCs communicate through a communications relay that hosts a publish/subscribe communications bus. This architecture removes the need for the DSMC or NLGCs to keep track of device addresses. Instead, each controller simply addresses the relay, which routes all communications between participating controllers.

At the fourth layer of control, the DSMC receives measurements of NLG net-load resources periodically (every 200 ms), and, on command, receives each NLGC's best net-load response options. The DSMC implements Mthd. 5 described in Part I to rapidly calculate the globally-optimal dispatch decision and then distributes dispatch commands to each NLGC. More specifics on the FFR optimal dispatch algorithm and its performance are given in Part I. The controllers for all four layers of the cosimulated NLGs are deployed on up to (20) 36-core linux servers.

III. SINGLE NET-LOAD GROUP EXPERIMENT RESULTS

A. Scenario

The aim of this scenario is to experimentally validate the response time of the FFR optimal dispatch solution when operating in single NLG dispatch mode and to highlight the performance and benefit of the NLU inverter optimal power point approach developed in Part I [1]. A single NLG with nine NLUs (NLU1–9 as described in Table IV), representing nine residential buildings with a total of 45 controlled net-load resources, is used as the basis for this scenario. Scenario data inputs include real-world appliance load profiles (as described in Section II-B2) as well as a variable solar irradiance profile (see Fig. 5). The physical appliance loads in NLU1–NLU2 are perturbed with user action throughout each trial. Multiple trials of the same scenario, but with different frequency event times and NLG participation factors, are conducted to demonstrate the performance under a variety of starting conditions. For all trials, the FFR dispatch event is initiated using an underfrequency event profile with a 60- to 59.9-Hz step. The NLU1 inverter is configured to monitor system frequency and to send a trigger to the NLGC to start its dispatch sequence whenever frequency declines to less than 59.92 Hz. NLU6 is configured to operate its inverter in optimal MPP tracking mode. That is, the inverter is operated with sufficient headroom to ensure the NLU is capable of providing an FFR event dispatch of any percentage of its controlled net-load capacity at all times, while minimizing the amount of headroom required (see Part

I for further details on the algorithm). To do so, NLGC #6 implements Algorithm 1 from Part I [1], uses the latest load measurements available, and sends power commands to the inverter every 1 s. The inverters of NLU1–5 and NLU7–8 are operated with a fixed headroom of 20%.

B. Results

1) *Summary*: Summary results are shown in Table II. It can be seen that the hierarchical approach was able to achieve the same global optimal dispatch solution as the mixed-integer linear program (MILP) solution in 23–56% less time. Global optimum dispatch results for comparison were obtained by running the MILP solution for the same set of net-load measurements and costs as were recorded by the hierarchical approach during its dispatch. Complete closed-loop responses, from system frequency change to all loads and inverters actuated, were achieved within 247 ms for all trials, well within the targeted 500 ms for FFR.

TABLE II.
SINGLE NLG EXPERIMENT RESULTS

Trial #		1	2	3	4
η		0.4	0.37	0.45	0.4
ΔP^* (W)	Dispatched	11,298	5,724	13,486	9,553
	Achieved	11,257	5,762	13,539	9,510
	(measured)				
	% diff.	-0.36%	0.67%	0.39%	-0.45%
Cost	Hierarchical Approach (HA)	24,108	10,710	24,455	16,517
	MILP solution	24,108	10,710	24,455	16,517
	% diff.	0%	0%	0%	0%
Solution Time (ms)	HA	88	90	79	85
	MILP	202	117	146	116
	% diff.	-56.4%	-23.2%	-45.8%	-26.9%
HA Solution Time (ms)	Detect Δf	55	57	64	63
	Dispatch alg.	88	90	79	85
	Misc. ¹	6	19	13	37
Breakdown (ms)	Initial response	149	166	156	185
	Complete response	225	208	231	247

2) *Single NLU Optimal MPP Tracking*: Here, detailed results are presented from Trial #2 and Trial #4 to highlight NLU optimal MPP tracking functionality. During both trials, two NLUs—NLU6 using optimal MPP tracking mode and NLU9 using regular MPP tracking mode—with identical equipment ratings, load profiles, and irradiance profile were run simultaneously. Fig. 5 shows simulation (expected operation, as predicted by a simulation using the same solar irradiance and load profile inputs) and experimental results for the two NLUs during Trial #2. It was verified that the NLU inverter operating in optimal MPP tracking mode operated off of the maximum power profile in proportion to the real-time calculated headroom, α (i.e., the “optMPP-meas” trace in Fig. 5 follows the simulated “optMPP-Sim” trace, which is the difference between “reg-Meas” and α). Measured operation matches the expected operation except at low power values, where the measured inverter maintains a minimum power of 300 W to avoid going into standby mode. Trial #2 did involve an FFR dispatch event, though neither NLU was dispatched.

In Trial #4, NLU6 is dispatched near $t = 6.1$ min by -1313 W, or 25.9% of its pre-event net-load consumption, as shown in Table III. The NLU6 dispatch set point was achieved entirely by the inverter. Had the NLU6 inverter been operating

¹Time for NLGC to send dispatch commands, communications delays, local controller processing time, and local device actuation.




Fig. 5. Trial #2 optMPP results: (top) Power output for inverter in optimal MPP tracking mode (“optMPP”) versus regular mpp tracking mode (“reg”). Measured experimental data (“Meas”) and predicted simulation results (“sim”) are compared. (middle) The real-time calculated α . (bottom) Load profiles for the five NLU loads.

in regular MPP tracking mode, the NLU would be limited to responding with a few discrete options. E.g., for $\eta = 25.9\%$, the two nearest discrete options would have been either 2.7% or 37.8% of the NLU’s net-load power. This demonstrates that the optimal MPP tracking strategy was effective for ensuring the NLU could be dispatched at any percentage of its net-load consumption. Detailed results for NLU6 and NLU9 operation during Trial #4 are shown in Fig. 6. Note that after the inverter is dispatched near $t = 6.1$ min, its output remains fixed at the minimum of the dispatched set point or the maximum available solar power.




Fig. 6. Trial #4 optMPP results

TABLE III.
NLU6 DISPATCH SCENARIOS

	Cost	Load Capacity (W) Under Mode			
		Opt. MPP		Reg. MPP	
		Before	After	Before	After #1
Inverter	4	-566	-1879	-1879	-1879
Load1	3	4110	4110	4110	4110
Load2	4	0	0	0	0
Load3	5	102	102	102	0
Load4	6	0	0	0	0
Load5	7	1420	1420	1420	0
Net		5066	3753	3753	2333
α^*		1294			
α		1313		0	
ΔP			-1313		-1420
η			0.259		0.027
Total cost			5252		9940

3) *NLG Optimal FFR Dispatch*: Measured RMS power results showing the operation of all NLUs and their constituent inverters and loads throughout Trial #4 are shown in Fig. 7,

and summary results are given in Table II. A dispatch event with participation factor $\eta = 0.4$ is completed, and assets in NLU1 (three loads and inverter), NLU2 (three loads and inverter), and NLU6 (inverter only) are dispatched to achieve a total NLG dispatch of $\Delta P_{meas}^* = 9510$ W. Instantaneous current and voltage waveforms of NLU1 and NLU2 assets close to the instant of the dispatch event are shown in Fig. 8. This plot confirms that initial response is seen approximately 185 ms after a step change in the system frequency. These results experimentally validate that an NLG of nine NLUs can be optimally dispatched to provide FFR within 247 ms and while matching the global optimum dispatch result.




Fig. 7. Trial #4 NLG dispatch results: RMS active power for the (a) three dispatched NLUs, (b) six nondispatched NLUs, (c) constituent loads of NLU2, (d) NLU inverters, and (e) the total NLG (all NLUs)

IV. FULL DISTRIBUTION FEEDER EXPERIMENT RESULTS

A. Scenario

This scenario experimentally validates the response time of the FFR optimal dispatch solution when operating in global dispatch mode with increasing numbers of NLGs controlled. The full distribution network configuration described in Section II-B1 and detailed in Table IV is used for this scenario. As before, scenario data inputs include real-world appliance load profiles as well as a variable solar irradiance profile, and the physical appliance loads in NLU1–NLU2 are perturbed with user action throughout each trial. Multiple trials of the same scenario, but with different frequency event times, numbers of NLGs being controlled, and total feeder response magnitude are conducted. For all trials, the FFR dispatch event is initiated using an underfrequency event profile with a 60- to 59.9-Hz step. The NLU1 inverter is configured to monitor system frequency and to send a trigger to the DSMC to start its dispatch sequence whenever the frequency declines to less




Fig. 8. Trial #4 NLG dispatch results: Instantaneous waveforms (top) NLU1 inverter current (“IL1”), NLU2 inverter current (“MISMA”), grid voltage (“VG”) and approximate grid frequency (“FREQ”). “FTRIG” goes from low to high when the local inverter controller detects the frequency event. “PTRIG” goes from high to low when the local inverter controller receives its dispatch set point from the NLGC. (bottom) Four NLU1 constituent load currents (“FRIDGE”, “WASHER”, “RANGE”, and “AIRCON”), the total load current of NLU1 (“HM3”), and the total load current of NLU2 (“HM2”).

TABLE IV.
EXPERIMENT CONFIGURATION SUMMARY

PHIL				Inverters		Load Type ²	Total Controllable Devices			
NLG	NLU	Node	HG	Type	Capacity (kW)		Physical		Virtual	
				Phys.	Sim.		PV	Batt.	Load	Inv.
1	1	1	1	T4	2	App	1	-	4	-
1	2	1	2	T1	2	App	1	-	4	-
1	3-8	2-7	3-8	T1	3	VLB	6	-	6	-
1	9	-	-	V	-	3	-	-	-	1
1	10	3	4	T2	3.6	3.6	1	-	-	-
2	11-14	4-7	5-8	T2	3.6	3.6	4	-	-	-
2	15-20	2-7	3-8	T5	3	3	-	6	-	-
3	21-26	2-7	3-8	T3	3.8	3.8	7 ²	-	-	-
3	27-30	-	-	V	-	P_R	-	-	-	4
4	31	8	9	T6	90	900	1	1	-	-
4	32	9	10	T6	90	90	1	1	-	-
5-26	10 ea.	-	-	V	-	P_R	-	-	-	220
27-114	10 ea.	-	-	V	-	P_R	-	-	-	880
115-224	10 ea.	-	-	V	-	P_R	-	-	-	1,100
Total Device Counts							87	8	8	6
								2,205	8,912	

than 59.92 Hz. All NLU inverters are operated with a fixed headroom of 20% for this scenario.

B. Results

1) *Summary*: Summary results for the full-scale experiments are shown in Table V. Four trials, each considering the same >100 physical HUT and increasing numbers of co-simulated DER, were conducted. Each trial also considered unique participation factors to demonstrate applicability over a range of participation levels. The hierarchical approach was able to achieve nearly the same (within 0.2%) global optimal dispatch solution as the MILP solution and does so in 41–83% less time. Total response time is crucial as the problem size (i.e., N_{nlg}) grows and should remain less than

²Appliance (App), virtual loads with load bank output (VLB), and virtual loads with simulated output (V) are used.

³Groups of 12 microinverters are individually controlled, but they have a combined capacity as listed.

⁴ P_R refers to an inverter capacity chosen at random to follow the statistical distribution of PV system capacities already in the field.

500 ms to show viability for providing FFR. The results in Table V demonstrate that the hierarchical approach developed here has a nearly constant dispatch solution time, whereas the MILP approach's solution time grows significantly with problem size. In particular, the MILP approach solution time for Trial #3 and Trial #4 is too long to be able to support FFR requirements, whereas the hierarchical approach is suitably fast. These dispatch solution time results confirm the conclusions of the simulation results comparing the two approaches in Part I. The hierarchical approach's dispatch solution times are nearly independent of problem size; however, note that the "Total gather all NL response options"⁵ and "Misc."⁶ times do increase some with problem size. This is a consequence of the DSMC needing to receive, process, and send out results from/to more NLGCs. The implementation of the DSMC used here completes the tasks of receiving data and sending out dispatch commands in a single thread (in series); however, the implementation can be upgraded to parallelize these tasks, which will allow the implementation to continue scaling as expected to upwards of $N_{nl} = 800$. The important result here is that the key parts of the dispatch solution—"Dispatch solution time" and "Get local NL response options (max)"—remain nearly constant despite greatly increasing problem size, demonstrating that the underlying algorithmic implementation is sound. For the largest problem tested, a case with $N_{nl} = 224$ that represents more than 10,000 net-load resources to be dispatched, the hierarchical approach is able to compute the entire dispatch and achieve output response from hardware within 500 ms.

TABLE V.
FULL FEEDER EXPERIMENT RESULTS

Trial #		1	2	3	4
η		0.05	0.13	0.20	0.25
N_{nl}		26	26	114	224
ΔP^* (W)	Dispatched - Total	172,359	482,646	854,693	1,034,529
	Dispatched - Com.	-	383,130	402,270	391,231
	Achieved (measured)	172,090	478,850	853,922	1,030,202
	% diff.	-0.26%	-0.79%	-0.09%	-0.42%
Cost	Hierarchical Approach (HA)	597,417	1,053,503	1,742,412	2,858,176
	MILP solution	597,218	1,051,374	1,740,908	2,857,474
	% diff.	0.03%	0.20%	0.09%	0.02%
Dispatch Solution Time (ms)	HA	187	197	195	205
	MILP	319	355	543	1,206
	% diff.	-41%	-45%	-64%	-83%
HA Solution Breakdown (ms)	Detect Δf	39	41	24	21
	Get local NL options (avg)	132	113	125	118
	Get local NL options (max)	149	147	151	163
	Total gather all NL response options ⁵	205	214	252	283
	Central dispatch alg.	38	50	44	42
	Misc. ⁶	52	46	61	79
	Total response - fastest hardware	334	351	381	425
	Total response - fastest COTS hardware	382	404	429	463

2) *Distribution Feeder-Scale Optimal FFR Dispatch:* Detailed results for the FFR dispatch of Trial #4 are presented

⁵Total time required to send out request for NL options, wait for responses, receive all responses, and concatenate them together.

⁶Time for send/receive dispatch commands and communications delays at central+local controllers, local controller processing, and local device actuation.

here. Fig. 9 shows that the majority of inverters in the experiment respond within 500 ms of the frequency event occurring. A key exception are the Type 3 inverters, which all respond slower than 500 ms and with some response time variation. This demonstrates the reality of managing a variety of device types in the field; some existing devices might not be fast enough to participate in FFR without further firmware upgrades. Note that the active power responses shown in Fig. 9 are based on RMS values that are calculated using a rolling window of multiple cycles. This means that the initial response times (i.e., when the slope of the trace changes) indicated are accurate, but the time to ramp up might be delayed versus reality. Thus, instantaneous current waveform measurements of selected inverters (one of each type) and the appliance loads are presented in Fig. 10 to confirm total response times. As shown, the appliance loads are all deferred quickly and the custom-designed Type 4 inverter also responds rapidly. All other inverters, with the exception of the Type 3 inverters, begin responding within 500 ms and generally ramp quickly. Fig. 11 shows additional results from a sampling of other NLGs which confirms that substation power ramps quickly and that the desired dispatch target is met within ~500 ms. The bottom graph of Fig. 11 shows that this rapid provision of FFR does have a small effect on voltages across the distribution feeder, though not enough to cause any deviations anywhere close to ANSI C84.1 [18] +/- 5% limits, especially for this relatively-stiff feeder.




Fig. 9. Full feeder results: RMS active power response for the more than 27 different hardware inverters in NLU1-NLU32

V. CONCLUSIONS

This article implemented a solution for FFR from optimally-coordinated net-load resources in a large-scale PHIL experiment representative of a real-world distribution network




Fig. 10. Full feeder results: oscilloscope view

model. The scalability and performance advantages of the developed hierarchical coordination algorithm shown via simulations in Part I [1] were confirmed via multiple experiments here. We successfully demonstrated that the approach developed can optimally coordinate more than 10,000 net-load resources across a distribution network and achieve hardware response within 500 ms, a feat not possible using state-of-the-art MILP approaches. This demonstration validates that coordination of behind-the-meter resources in distribution networks can be achieved rapidly enough and on a sufficiently large scale to provide valuable FFR services to the bulk power system. Although additional implementation work to achieve robust, large-scale field deployments and further development of regulatory and market structures will be needed to allow full adoption of this scheme, this work provides a first-of-its-kind experimental demonstration that this approach is viable.

REFERENCES

- [1] B. Lundstrom, S. Patel, and M. V. Salapaka, "Distribution Feeder-Scale Fast Frequency Response via Optimal Coordination of Net-load Resources-Part I: Solution Design," *IEEE Transactions on Smart Grid*.
- [2] W. Ren, M. Steurer, and L. Qi, "Evaluating dynamic performance of modern electric drives via power-hardware-in-the-loop simulation," in *2008 IEEE International Symposium on Industrial Electronics*, Jun. 2008, pp. 2201–2206, iSSN: 2163-5145.
- [3] B. Lundstrom, B. Mather, M. Shirazi, and M. Coddington, "Methods and Implementation of Advanced Unintentional Islanding Testing Using Power Hardware-in-the-Loop (PHIL)," in *IEEE Photovoltaic Specialists Conference (PVSC)*, Tampa Bay, FL, 2013.
- [4] B. Lundstrom, S. Chakraborty, G. Lauss, R. Bründlinger, and R. Conklin, "Evaluation of system-integrated smart grid devices using software- and hardware-in-the-loop," in *2016 IEEE Power Energy Society Innovative*




Fig. 11. Full feeder results: RMS active power response for the appliance loads in NLU2, a selection of the 224 NLGs, all residential and commercial NLUs, and at the distribution feeder head at the substation. Voltages from a selection of points across the full distance from substation are also shown.

Smart Grid Technologies Conference (ISGT), Sep. 2016, pp. 1–5, iSSN: 2472-8152.

- [5] P. C. Kotsampopoulos, F. Lehuss, G. F. Lauss, B. Bletterie, and N. D. Hatziaargyriou, "The Limitations of Digital Simulation and the Advantages of PHIL Testing in Studying Distributed Generation Provision of Ancillary Services," *IEEE Transactions on Industrial Electronics*, vol. 62, no. 9, pp. 5502–5515, Sep. 2015.
- [6] X. Xi, H. Geng, G. Yang, Y. Zhou, J. Lin, and Y. Song, "Small-signal stability analysis and PHIL experiment validation of power system with WECS voltage/frequency ancillary control," in *IECON 2015 - 41st Annual Conference of the IEEE Industrial Electronics Society*, Nov. 2015, pp. 002 544–002 549.
- [7] O. Tremblay, H. Fortin-Blanchette, R. Gagnon, and Y. Brissette, "Contribution to stability analysis of power hardware-in-the-loop simulators," *IET Generation, Transmission & Distribution*, vol. 11, no. 12, 2017.
- [8] S. Chakraborty, A. Hoke, and B. Lundstrom, "Evaluation of multiple inverter volt-Var control interactions with realistic grid impedances," in *Power & Energy Society General Meeting, 2015 IEEE*. IEEE, 2015, pp. 1–5. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=7285795
- [9] A. F. Hoke, A. Nelson, S. Chakraborty, F. Bell, and M. McCarty, "An Islanding Detection Test Platform for Multi-Inverter Islands Using Power HIL," *IEEE Transactions on Industrial Electronics*, vol. 65, no. 10, pp. 7944–7953, Oct. 2018.
- [10] G. Lauss, M. O. Faruque, K. Schoder, C. Dufour, A. Viehweider, and J. Langston, "Characteristics and Design of Power Hardware-in-the-Loop Simulations for Electrical Power Systems," *IEEE Transactions on Industrial Electronics*, vol. 63, no. 1, pp. 406–417, Jan. 2016.
- [11] B. Palmintier, B. Lundstrom, S. Chakraborty, T. Williams, K. Schneider, and D. Chassin, "A Power Hardware-in-the-Loop Platform With Remote Distribution Circuit Cosimulation," *IEEE Transactions on Industrial Electronics*, vol. 62, no. 4, pp. 2236–2245, Apr. 2015.
- [12] B. Lundstrom, S. Patel, S. Attree, and M. V. Salapaka, "Fast Primary Frequency Response using Coordinated DER and Flexible Loads: Framework and Residential-scale Demonstration," in *IEEE Power and Energy Society General Meeting*, Aug. 2018.
- [13] B. Palmintier, B. Lundstrom, S. Chakraborty, T. Williams, K. Schneider, and D. Chassin, "A Power Hardware-in-the-Loop Platform With Remote Distribution Circuit Cosimulation," *IEEE Transactions on Industrial*

Electronics, vol. 62, no. 4, pp. 2236–2245, Apr. 2015.

- [14] Opal-RT Technologies, “ePHASORSIM.” [Online]. Available: <https://opal-rt.com/systems-phasorsim/>
- [15] J. Z. Kolter and M. J. Johnson, “REDD: A public data set for energy disaggregation research,” in *SustKDD workshop on Data Mining Applications in Sustainability*, 2011.
- [16] Raspberry Pi Foundation, “Raspberry Pi 3 Model B+.” [Online]. Available: <https://raspberrypi.org/products/raspberry-pi-3-model-b-plus/>
- [17] iMatix, “ZeroMQ Distributed Messaging.” [Online]. Available: <http://zeromq.org/>
- [18] *American National Standard for Electric Power Systems and Equipment—Voltage Ratings (60 Hz)*. ANSI Standard C84.1, 2016.



Blake Lundstrom (M’11–SM’19) received the B.S. and M.S. degrees in electrical engineering from the Colorado School of Mines, Golden, CO, USA in 2010 and 2013, respectively, and the Ph.D. degree in electrical engineering from the University of Minnesota, Minneapolis, MN, USA in 2020. He is a Senior Research Engineer in the Power System Engineering Center at the National Renewable Energy Laboratory, Golden, CO, USA, where he has been employed since 2009. His research involves design of control techniques for managing large-scale

deployments of distributed energy resources, distribution system modeling, power electronics controls development, and the use and development of real-time hardware-in-the-loop techniques for energy systems integration research. Dr. Lundstrom is a Vice Chair for IEEE P2004: *Recommended Practice for Hardware-in-the-Loop Simulation Based Testing of Electric Power Apparatus and Controls* and is a licensed Professional Engineer in the State of Colorado.



Sourav Patel (S’18) received his B.Tech. degree in Instrumentation and Control Engineering from the National Institute of Technology, Jalandhar, India, in 2011. In the same year, he joined National Thermal Power Corporation Ltd. in Kaniha, Odisha, India as a Control and Instrumentation Engineer. He received his M.S. degree in Electrical Engineering from the University of Minnesota, Minneapolis, MN, USA, in 2018 where currently, he is working towards his Ph.D. degree in Electrical Engineering. His research

interests include control and systems theory, coordination and communication protocols for Distributed Energy Resources towards smart grid applications.



Murti V. Salapaka (F’19) received the B.Tech. degree in Mechanical Engineering from the Indian Institute of Technology, Madras, in 1991 and the M.S. and Ph.D. degrees in Mechanical Engineering from the University of California at Santa Barbara, in 1993 and 1997, respectively. He was a faculty member in the Electrical and Computer Engineering Department, Iowa State University, Ames, IA, USA, from 1997 to 2007. Currently, he is the Director of Graduate Studies and the Vincentine Hermes Luh Chair Professor in the Electrical and Computer

Engineering Department, University of Minnesota, Minneapolis, MN, USA. His research interests include control and network science, nanoscience and single molecule physics. Dr. Salapaka received the 1997 National Science Foundation CAREER Award and is an IEEE fellow.