Frequency Regulation with Heterogeneous Energy Resources: A Realization using Distributed Control

Tor Anderson* Manasa Muralidharan* Priyank Srivastava* Hamed Valizadeh Haghi
Jorge Cortés Jan Kleissl Sonia Martínez Byron Washom

***Abstract*—This paper presents one of the first real-life demonstrations of coordinated and distributed resource control for secondary frequency response in a power distribution grid. A series of tests involved up to 69 heterogeneous active distributed energy resources consisting of air handling units, unidirectional and bidirectional electric vehicle charging stations, a battery energy storage system, and 107 passive distributed energy resources consisting of building loads and solar photovoltaic systems. The distributed control setup consists of a set of Raspberry Pi endpoints exchanging messages via an ethernet switch. Actuation commands for the distributed energy resources are obtained by solving a power allocation problem at every regulation instant using distributed ratio-consensus, primal-dual, and Newton-like algorithms. The problem formulation minimizes the sum of distributed energy resource costs while tracking the aggregate setpoint provided by the system operator. We demonstrate accurate and fast real-time distributed computation of the optimization solution and effective tracking of the regulation signal over 40 min time horizons. An economic benefit analysis confirms eligibility to participate in an ancillary services market and demonstrates up to $53k of potential annual revenue for the selected population of distributed energy resources.**

I. Introduction

Many recent efforts seek to integrate renewable energy resources with the power grid to reduce the carbon footprint. The high variability associated with wind and solar power can be balanced using distributed energy resources (DERs) providing ancillary services such as frequency regulation. Consequently, there is a growing interest among market operators in DER aggregations with flexible generation and load capabilities to balance fluctuations in grid frequency and minimize area control errors (ACE). The fast ramping rate and minimal marginal standby cost put many DERs at an advantage against conventional generators and make them suitable for participation in the frequency regulation market.

The fast ramping rates reduce the required power capacity of DERs to only 10% of an equivalent generator to balance a frequency drop within 30 s [1]. However, most individual DERs have small capacities, typically on the order of kWs compared to 10 s of MW for conventional frequency control resources. Commanding the required thousands to millions of DERs to replace existing frequency regulation resources over a large balancing area entails aggregating DERs that are distributed at end points all over the grid on customer premises. The dynamic nature, large number, and distributed location of DERs requires coordination. This is in contrast to existing frequency regulation [2] implementation with conventional energy resources. For example, California Independent System Operator (CAISO) requires all generators to submit their bids once per regulation interval. Then, the setpoints are assigned centrally to all resources every 2-4 s without any consideration of operational costs [3]. While distributed control has the potential to enable DER participation in the frequency regulation market (e.g., [4]), there is a general lack of large-scale testing to prove its effectiveness for widespread adoption by system operators. The 2017 National Renewable Energy Laboratory Workshop on Autonomous Energy Grids [5] concluded that "A major limitation in developing new technologies for autonomous energy systems is that there are no large-scale test cases (...). These test cases serve a critical role in the development, validation, and dissemination of new algorithms".

The results of this paper are the outcome of a project under the ARPA-e Network Optimized Distributed Energy Systems (NODES) program[1], which postulates DER aggregations as virtual power plants that enable variable renewable penetrations of at least 50%. The vision of the NODES program was to employ state-of-the-art tools from control systems, computer science, and distributed systems to optimally respond to dynamic changes in the grid by leveraging DERs while maintaining customer quality of service. The NODES program required testing with at least 100 DERs at power. Here, we demonstrate the challenges and opportunities of testing on a heterogeneous fleet of DERs for eventual operationalization of optimal distributed control at frequency regulation time scales.

Literature Review. To the best of our knowledge, real-world testing of frequency regulation by DERs has been limited. A Vehicle-to-Grid (V2G) electric vehicle (EV) [6] and two Battery Energy Storage System (BESS) [7] provided frequency regulation. 76 bitumen tanks were integrated with a simplified power system model to provide frequency regulation via a decentralized control algorithm in [8]. In buildings, a decentralized control algorithm controlled lighting loads in a test

Tor Anderson, Priyank Srivastava, Jorge Cortés, and Sonia Martínez are with the Department of Mechanical and Aerospace Engineering, UC San Diego, CA, USA. Email: {tka001, psrivast, cortes, soniamd}@ucsd.edu.

Manasa Muralidharan, Jan Kleissl, and Hamed Valizadeh Haghi are with the Department of Mechanical and Aerospace Engineering and the Center for Energy Research, UC San Diego, CA, USA. Email: m1murali@ucsd.edu, valizadeh@ieee.org, jkleissl@ucsd.edu.

Byron Washom is with Strategic Energy Initiatives, UC San Diego, CA, USA. Email: bwashom@ucsd.edu.

This research was supported by the Advanced Research Projects Agency - Energy (ARPA-e) under the NODES program, Cooperative Agreement DE-AR0000695.

*: These authors contributed to the paper equally.

[1]https://arpa-e.energy.gov/arpa-e-programs/nodes

room [9], centralized frequency control was applied to an air handling unit (AHU) [10], [11], an inverter and four household appliances [12], and four heaters in different rooms [13]. A laboratory home with an EV and an AHU, and a number of simulated homes were considered for demand response in [14] through an aggregator at a 10 s level. Technologies for widespread, but centrally controlled, cycling of air conditioners directly by utilities cf. [15] and aggregators are common place for peak shifting, but occur over time scales of minutes to hours. Industrial solutions enabling heterogeneous DERs to track power signals also exist, but they are either centralized, cf. [16] or require all-to-all communication [17].

Our literature review exposes the following limitations: (i) centralized control or need for all-to-all communication [6], [7], [10]–[17], which does not scale to millions of DERs; (ii) small numbers of DERs [6], [7], [10]–[14]; (iii) lack of diversity in DERs [6]–[11], [13], with associated differences in tracking time scales and accuracy. No trial has been reported that demonstrated generalizability to a real scenario with (i) scalable distributed control and a (ii) large number of (iii) heterogeneous DERs.

Statement of Contributions. To advance the field of real-world testing of DERs for frequency control, we conduct a series of tests using a group of up to 69 active and 107 passive heterogeneous DERs on the University of California, San Diego (UCSD) microgrid [18]. To the best of the authors' knowledge, this is the first work to consider such a large, diverse portfolio of real physical DERs for secondary frequency response. As such, the major contributions of this work are:

- A detailed account of the testbed, including the DER actuation and sampling interfaces, the distributed optimization setup, and communication framework.
- A description of techniques to work around technical barriers, provision of lessons learned, and suggestions for future improvement.
- Evaluation of the performance of both the cyber and physical layers, including an evaluation of eligibility requirements for and the economic benefit of participating in the ancillary services market.

Paper Overview. Frequency regulation is simulated on the UCSD microgrid using real controllable DERs (Section III-C) to follow the Pennsylvania-New Jersey-Maryland Interconnection (PJM) RegD signal [19] interpolated from 0.5Hz to 1Hz (Sections III-B). The DER setpoint tracking is formulated as a power allocation problem at every regulation instant (Section III-A), and uses three types of provably convergent distributed algorithms from [20]–[23] to solve the optimization problem; see Appendix A of [24], removed from this version for brevity. Setpoints are computed distributively on multiple Raspberry Pi's communicating via ethernet switches (Section III-D). The setpoints are implemented on up to 176 DERs at power using dedicated command interfaces via TCP/IP communication (Section III-E), the DER power outputs monitored (Section III-F), and their tracking performance evaluated (Section III-G). Results (Section V) for the various test scenarios described in Section IV show that the test system tracks the signal with reasonable error despite delays in response and inaccurate tracking behavior of some groups of DERs, and qualifies for participation in the PJM ancillary services market .

II. Problem Setting

This paper validates real-world DER controllability for participation in secondary frequency regulation through demonstration tests implemented on a real distribution grid. The tests showcase the ability of aggregated DERs to function as a single market entity that responds to frequency regulation requests from the independent system operator (ISO) by optimally coordinating DERs. The goal is to monitor and actuate a set of real controllable DERs to collectively track a typical automatic generation control (AGC) signal issued by the ISO.

Three different distributed coordination schemes optimize the normalized contribution of each DER to the cumulative active power signal. Unlike simulated models, the use of real power hardware exposes implementation challenges associated with measurement noise, sampling errors, data communication problems, and DER response. To that end, precise load tracking is pursued at timescales that differ by DER type consistent with individual DER responsiveness and communication latencies, yet meet frequency regulation requirements in aggregation.

The 69 kV substation and 12 kV radial distribution system owned by UCSD to operate the 5 km^2 campus was the chosen demonstration testbed. It has diverse energy resources with real-time monitoring and control capabilities, allowing for active load tracking. This includes over 3 MW of solar photovoltaic (PV) systems, 2.5 MW/5 MWh of BESS, building heating ventilation and air conditioning (HVAC) systems in 14 million square feet of occupied space, and over 200 unidirectional V2G (V1G) [25] and V2G EV chargers. The demonstration tests used a representative population of up to 176 such heterogeneous DERs to investigate tracking behavior of specific DER types as well as their cooperative tracking abilities. While the available DER capacity at UCSD far exceeds the minimum requirements for an ancillary service provider set by most ISOs (typically $\sim$ 1 MW), logistical considerations and controller capabilities dictated the choice of a DER population size with less aggregate power capacity (up to 184 kW) for this demonstration. Since this magnitude of power is insufficient to measurably impact the actual grid frequency, we chose to simulate frequency regulation by following a frequency regulation signal.

III. Test Elements

Here, we elaborate on the different elements of the validation tests. These include the optimization formulation employed to compute DER setpoints (Section III-A), the reference AGC signal (Section III-B) and types of DERs used to track it (Section III-C), the computing platform (Section III-D), the actuation (Section III-E) and monitoring interfaces (Section III-F), the performance metrics used to assess the cyber and physical layers, and eligibility for market participation (Section III-G).

A. Optimization Formulation

The optimization model for AGC signal tracking using DERs can be mathematically stated as a separable resource allocation problem subject to box constraints as follows:

$$
\begin{aligned}
\min_{p\in\mathbb{R}^n} \; & f(p) = \sum_{i=1}^{n} f_i(p_i), \\
\text{s.t.} \; & \sum_{i=1}^{n} p_i = P_{\text{ref}}, \\
& p_i \in [\underline{p}_i, \overline{p}_i], \quad \forall i \in \mathcal{N} = \{1,\dots,n\}.
\end{aligned} \tag{1}
$$

The agents $i \in \mathcal{N}$ each have local ownership of a decision variable $p_i \in \mathbb{R}$, representing an active power generation or consumption quantity (setpoint), a local convex cost function f_i, and local box constraints $[\underline{p}, \overline{p}]$, representing active power capacity limits. P_{ref} is a given active power reference value determined by the ISO and transmitted to a subset of the agents as problem data, see e.g. [26]. P_{ref} is a signal that changes over time, so a new instance of (1) is solved in real-time 1 s intervals corresponding to these changes. Note that with just 1 s difference between the instances, the box constraints might also change due to the limited ramp rates of DERs. In this work we consider them constant and assume (1) is feasible.

For the validation tests, we used two types of cost functions: constant and quadratic. Constant functions were used for the Ratio-Consensus (RC) solver [20], which turns the optimization into a feasibility problem. Quadratic functions were used for the primal-dual based (PD) [21], [22] and Distributed Approximate Newton Algorithm (DANA) [23] methods. In short, RC prescribes dynamics which seek to achieve *consensus* on a *ratio* of operating capacity with respect to $\underline{p}_i, \overline{p}_i$ so that the agents achieve $\sum_i p_i = P_{\text{ref}}$. PD and DANA each are Lagrangian-based dynamics; in particular, PD is gradient-based ("first-order") and DANA is Newton-based ("second-order"). See Appendix A of the extended version [24] for more technical detail on these algorithms. The quadratic functions were artificially chosen to produce satisfactorily diverse and representative solutions to (1) for each DER population. Costs associated with a physical or economic metric (e.g. deviation from a building setpoint for AHUs, user-specified charging demands for V1G and V2Gs, and resistive losses in a BESS) are of great interest, but are far from trivial to model and thus not the focus of this study. We split the total time period of the signal, P_{ref} into three equal segments, and implemented RC, PD, and DANA in that order. Box constraints $[\underline{p}_i, \overline{p}_i]$ are given in Table I and were centered at zero for simplicity; for example, an AHU i with 2 kW capacity has $[\underline{p}_i, \overline{p}_i] = [-1, 1]$, while a V2G j with ± 5 kW capacity has $[\underline{p}_j, \overline{p}_j] = [-5, 5]$.

B. Regulation Signal

The 40 min RegD signal published by PJM [19] served as the reference AGC signal for the validation tests, and was used to obtain the value for P_{ref} in (1). The normalized RegD signal, contained in $[-1, 1]$ (see Figure 1), was interpolated from 0.5 Hz to 1 Hz. The signal was then treated by subtracting the normalized contributions of building loads and PV systems, cf. Section III-C. Finally, the normalized signal was scaled by a factor proportional to the total DER capacity $\sum_i (\overline{p}_i - \underline{p}_i)$ before sending to the optimization solvers. More precisely,

$$
P_{\text{ref}} = \beta \frac{\sum_i (\overline{p}_i - \underline{p}_i)}{\| P_{\text{RegD}} + P_{\text{PV}} - P_{\text{b}} \|_\infty} \left(P_{\text{RegD}} + P_{\text{PV}} - P_{\text{b}} \right), \tag{2}
$$

where P_{RegD} refers to the normalized RegD signal data, P_{PV} and P_{b} respectively refer to the normalized PV generation and building load data obtained from the UCSD ION server as described in Section III-F, and $0 < \beta < 1$ is an arbitrary scaling constant. Note that this results in a different target signal P_{ref} for the different test scenarios considered in Section IV due to the different power ratings of the DERs (cf. Section III-C) used across the tests. For most test scenarios, $\beta = 0.75$ to prevent extreme set points that would require all DERs to operate at either $\overline{p}_i$ or $\underline{p}_i$ simultaneously, which may be infeasible in some time steps due to slower signal update times, see Table II. Each P in (2) is a vector with 2401 elements corresponding to each 1 s time step's instance of (1) over the 40 min time horizon.

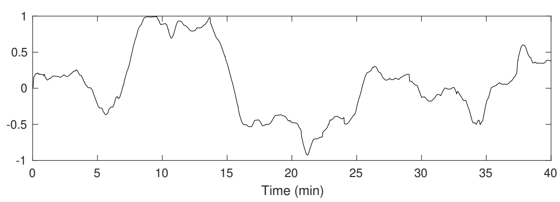

Fig. 1: Normalized PJM RegD signal.

C. DERs

The reference AGC signal was to be collectively tracked using DERs consisting of HVAC AHUs, BESS, V1G and V2G EVs, PV systems, and whole-building loads. Since PV systems and (non-AHU) building loads were not controllable, they participated in the test as passive DERs. Consequently, the active DERs were commanded to track a modified target signal derived by subtracting the net active power output of passive DERs from the reference AGC signal and applying appropriate scaling (cf. Section III-B). Table I lists the typical net power capacity $\overline{p}_i - \underline{p}_i$ of the different active DER types.

TABLE I: Typical power rating of active DER types

DER Type	AHU	V1G EV	V2G EV	BESS
Typical power rating per DER type	2 kW	3.3 kW (Tests 0 & 1), 4.9 kW (Test 2)	$\pm$ 5 kW	$\pm$ 3 kW

The contribution of each active DER to the target signal was defined with respect to a baseline power, around which $[\underline{p}_i, \overline{p}_i]$ was centered, to enable tracking of both positive and negative ramps in the target signal. For DERs like V2G EVs and BESS, which were capable of power adjustments in both directions, the baseline was 0 kW. The baseline for V1G EVs was defined to be halfway between their allowed minimum and maximum charging rates, where the former was restricted by the SAE J1772 charging standard to 1.6 kW. Similarly, the baseline for AHUs was defined to be half of their power draw when on. Further, since AHUs were limited to binary on-off operational states, the continuous and arbitrarily precise AHU setpoints obtained by solving (1) were rounded to the closest discrete setpoint obtained from a combination of on-off states before actuation.

AHU control was restricted, by UCSD Facilities Management, to specifying only DER setpoints and duration of actuation; since building automation controllers could not be modified, model-based designs were impossible. This was to avoid malfunctioning or disruptions to real physical infrastructure in the networked building management system that also controls lighting, security, and fire protection systems.

D. Computing Setup

The DER active power setpoints were computed for the entire 40-min test horizon prior to any device actuation using a set of 9 Linux-based nodes. The nodes C1-C9 communicate with each other over an undirected ring topology, cf. Fig. 2. As one of the sparsest network topologies, where message passing occurs only between a small number of neighbors, the ring topology presents a challenging scenario for distributed control. Since there were more active DERs than computing nodes, the 9 nodes were mapped subjectively to the 69 active DERs such that nodes C1-C2 computed the actuation setpoints for the AHUs, C3 for V1G EVs, C4-C8 for V2G EVs and C9 for the BESS.

The computing steps are summarized in Algorithm 1. Each computing node generated actuation commands as CSV files containing the power setpoints for their respective group of DERs at a uniform update rate of 1 Hz. Preliminary testing revealed different response times across DER types, with AHUs and V1G EVs exhibiting slower response than other active DER types. DERs with response times greater than 1 s were subject to a stair-step control signal with a signal update time consistent with DER responsiveness and constant setpoints during intermediate time steps. Table II lists the signal update times for the different DER types.

Algorithm 1 Computing process

Require: Map $f : C_i \to$ DER-type
1: Initialize time of last solution update $t_{\texttt{sol-update}_i} = 0$, initial setpoints for DERs mapped to computing node C_i as $P_{f(C_i)}, \forall i \in \{1, \dots, 9\}$
2: **for** $k = 0, \dots, 2400$ **do**
3:     **for** $i = 1, \dots, 9$ **do**
4:         **if** $k - t_{\texttt{sol-update}_i} == t_{\texttt{signal-update}_i}$ **then**
5:             Solve (1) to update $P_{f(C_i)}(k)$
6:             $t_{\texttt{sol-update}_i} = k$
7:         **end if**
8:         $P_{f(C_i)}(k) \leftarrow P_{f(C_i)}(t_{\texttt{sol-update}_i})$
9:         **if** $\mathrm{mod}(k, 60) == 0$ **then**
10:             Send $P_{f(C_i)}(k)$ to DER type, $f(C_i)$
11:         **end if**
12:     **end for**
13: **end for**

E. Actuation Interfaces and Communication Framework

The actuation commands were issued using fixed IP computers through dedicated interfaces that varied by DER type as depicted in Fig. 2. The setpoints for AHUs were issued through a custom Visual Basic program that interfaced with the Johnson Control Metasys building automation software. The power rate of the BESS was set via API-based communication with a dedicated computer that controlled the battery inverter. The V1G and V2G EVs charging rates were adjusted through proprietary smart EV charging platforms of the charging station operators. EVs using ChargePoint® V1G stations were manually controlled via the load shedding feature of ChargePoints station management software. The actuation of EVs using PowerFlex® V1G chargers and Nuvve® V2G chargers was automated and commands were issued via API-based communication.

F. Power Measurements

The active power of all DERs was metered at a 1 Hz frequency. The power outputs of individual PV systems and building loads were obtained prior to the test from their respective ION meters by logging data from the UCSD ION Supervisory Control and Data Acquisition (SCADA) system and aggregated to obtain the total power output of all PVs and building loads. A moving average filter with a 20 s time horizon was used to remove noise from the aggregate measured data for these passive DERs. V2G EVs and BESS power data were acquired using the same interfaces that were used for their actuation, which logged data from dedicated power meters.

Since neither AHUs nor the ChargePoint V1G EVs had dedicated meters, they were monitored via their respective building ION meters by subtracting a baseline building load from the building meter power output. Assuming constant baseline building load, any change in the meter outputs can be attributed to the actuation of AHUs and V1G EVs. This assumption is justifiable considering the tests were conducted at 0400 PT to 0600 PT on a weekend, when building occupancy was likely zero and building load remained largely unchanged. Noise in the ION meter outputs observed as frequent 15 - 30 kW spikes in the measured data for AHUs (Fig. 3) and ChargePoint V1G EVs was treated by removing outliers and passing the resulting signal through a 4 s horizon moving average filter. Here, outliers refer to points that change in excess of 50% of the mean of the 40 min signal in a 1 s interval.

G. Performance Metrics

The performance of the distributed implementation (cyber-layer) was measured by the normalized mean-squared-error (MSE) between the distributed and true (i.e. exact) centralized optimization solutions. The true solutions were computed for each instance of (1) using a centralized CVX solver in MATLAB [27]. The MSE was normalized by dividing by the mean of the squares of the true solutions.

The tracking performance of the DERs was evaluated through (i) the root-mean-squared-error (RMSE) in tracking

$$\mathrm{RMSE} = \sqrt{\frac{\sum_{t=1}^{T}(P_t^{\mathrm{prov}} - P_t^{\mathrm{tar}})^2}{\sum_{t=1}^{T}(P_t^{\mathrm{tar}})^2}}, \tag{3}$$

where P_t^{prov} is the total power that was provided (measured), and P_t^{tar} is the target (commanded) regulation power at time

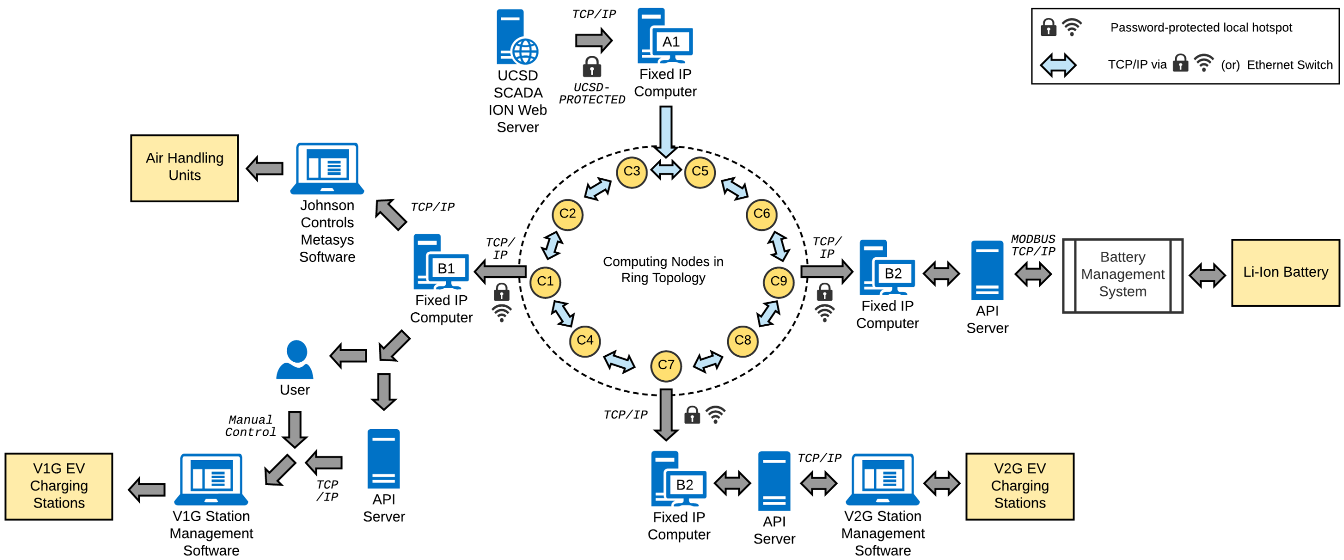

Fig. 2: Communication architecture for computation and actuation of control policies.

step $t \in \{1, \ldots, T = 2401\}$; and (ii) the tracking delay, computed as the time shift of the measured signal which yields the lowest RMSE between the commanded and measured signals. The sum of the delays due to local computation and communication between the computing nodes is capped by the algorithm computation time, and would be less than 1 s. Therefore, these delays are not explicitly considered in the tracking delay calculation, and the computed tracking delay only includes the device response times and measurement delays.

The PJM Performance Score S following [28, Section 4.5.6] was computed as a test for eligibility to participate in the ancillary services market, and is given by the mean of a Correlation Score S_c, Delay Score S_d, and Precision Score S_p:

$$S_c = \frac{1}{T-1}\sum_{t=1}^{T}\frac{(P_t^{\text{prov}} - \mu^{\text{prov}})(P_t^{\text{tar}} - \mu^{\text{tar}})}{\sigma^{\text{prov}}\sigma^{\text{tar}}},$$

$$S_d = \left|\frac{\delta - 5 \text{ min}}{5 \text{ min}}\right|, \quad S_p = 1 - \frac{1}{T}\sum_{t=1}^{T}\left|\frac{P_t^{\text{prov}} - P_t^{\text{tar}}}{\mu^{\text{tar}}}\right|,$$

$$S = 1/3(S_c + S_d + S_p),$$

where P_t^{prov} and P_t^{tar} are as in (3), $\mu^{\text{prov}}, \mu^{\text{tar}}$ and $\sigma^{\text{prov}}, \sigma^{\text{tar}}$ denote their respective means and standard deviations, and δ is the corresponding maximum delay in DER response for when S_c was maximized. A performance score of at least 0.75 is required for participating in the PJM ancillary services market.

IV. Test Scenarios

In this section, we describe the test scenarios carried out on the UCSD microgrid elaborating on the challenges we faced and the differences across the tests, summarized by type of DER in Table II.

A. Commonalities

A series of three tests were conducted on December 12, 2018 (Test 0), April 14, 2019 (Test 1) and December 17, 2019 (Test 2). All three tests involved a 40 min preparatory run followed by a 40 min final test. Table II lists the type of DERs across the tests. All tests were carried out during non-operational hours (between 0400 PT and 0540 PT) to avoid potential disruptions to building occupants with the exception of V1G EVs in Test 2, which were tested at the start of the work day (0900 - 1010 PT) to maximize fleet EV availability (cf. Section IV-D). Day-time PV output data from February 24, 2019 was used as a proxy for an actual daytime PV signal.

TABLE II: Characteristics of each test by DER type.

<table>
<tr><th>DER Type</th><th>AHU</th><th>V1G EV</th><th>V2G EV</th><th>BESS</th></tr>
<tr><td># DERs - Test 0</td><td>7</td><td>4</td><td>5</td><td>1</td></tr>
<tr><td># DERS - Test 1</td><td>34</td><td>29</td><td>5</td><td>1</td></tr>
<tr><td># DERs - Test 2</td><td>34</td><td>17</td><td>6</td><td>1</td></tr>
<tr><td>Signal updates</td><td>1 m</td><td>5 m (Tests 0 & 1), 1 m (Test 2)</td><td>1 s</td><td>20 s</td></tr>
<tr><td>DER Actuation</td><td>Synchronous (Tests 0 & 1), Two-stage: Stage 1 (Test 2)</td><td colspan="3">Synchronous (Tests 0 & 1), Two-stage: Stage 2 (Test 2)</td></tr>
<tr><td>Operation Mode</td><td>Automatic</td><td>Manual (Tests 0 & 1), Automatic (Test 2)</td><td colspan="2">Automatic</td></tr>
<tr><td>Time of test</td><td>0400 - 0500 PT</td><td colspan="3">0400 - 0500 PT (Tests 0 & 1), 0900 - 1010 PT (Test 2)</td></tr>
<tr><td>Computing setup</td><td colspan="4">Semi-centralized using ROS (Tests 0 & 1), Fully distributed using Raspberry Pi (Test 2)</td></tr>
</table>

B. Test 0

Test 0 was a preliminary calibration that was used to examine the response times and tracking behavior of every DER type and detect issues related to communication and actuation.

1) DERs: Test 0 used only a representative sample of 17 DERs. The V1G and V2G population was composed of UCSD fleet EVs plugged in at ChargePoint and Nuvve charging stations, respectively.

2) Computing Setup: 9 laptops running a Robotic Operating System (ROS) communicated via local Wi-Fi hotspot to implement the distributed coordination algorithms and compute the DER setpoints.

3) Actuation: All DERs were actuated synchronously.

C. Test 1

Test 1 was identical to Test 0 except in the number of DERs utilized.

1) DERs: Test 1 used a larger population of 69 active DERs and 107 passive DERs.

2) Computing Setup: The same semi-centralized ROS-based computing setup as in Test 0 was used in Test 1. Given that the available power capacity of fast-responding DERs such as V2G and BESS was smaller than slow-responding DERs, the steep ramping demands of the target signal were met by upscaling the power of the fast responding DERs in solving for the contribution of individual DERs. Another option would have been to reduce the number of slow responding DERs, but the funding agency stipulated prioritizing the number and types of heterogeneous DERs over accuracy in signal tracking. A real DER aggregator would instead require a more balanced capacity of slow and fast DERs to ensure feasibility of tracking these ramp features.

3) Actuation: All DERs were actuated synchronously. Since the ChargePoint V1G EVs in Test 1 were operated via manual input of DER setpoints (an interface to their API had not been developed yet), to avoid overloading the (human) operators, they were grouped into three groups and actuated in a staggered fashion such that each of the three groups maintained a signal update time of 5 min but were commanded 1 min apart from each other.

D. Test 2

Test 2 also used the entire population of DERs but substituted the cumbersome V1G population with more capable V1G chargers and used a new distributed computing setup and method of actuation based on lessons learned from Test 1.

1) DERs: The V1G EVs used in Test 1 performed poorly owing to an unreliable actuation-interface that experienced seemingly random stalling and lacked automated control capabilities. Therefore, 17 PowerFlex V1G charging stations at one location replaced the distributed 29 V1G charging stations used in Test 1. Since the PowerFlex interface did not permit actuating individual stations, the 17 charging stations participated in the test as a single aggregate DER. The 0930 1010 PT timing of the V1G EV part of the test coincided with the start of the workday and a V1G EV population that had only recently plugged in and therefore had ample remaining charging capacity. The EVs were contributed by UCSD employees and visitors randomly plugging in at the PowerFlex charging stations just before the start of the trial. An aggregate signal of 15 kW to 19 kW was distributed equally amongst the 17 EVs.

In addition to the new V1G EVs, the V2G population in Test 2 was replaced with a different set of Nuvve chargers to resolve a tracking/noise issue during discharge-to-grid observed in Test 1 and expanded to include an additional charger.

2) Computing Setup: Test 2 featured a fully distributed architecture that consisted of a network of Raspberry Pis that asynchronously communicated with each other via an ethernet switch. In addition, a modified synchronization technique was implemented in the software which improved the fidelity and robustness of message-passing. This upgraded message-passing framework and synchronization technique for both software and hardware resulted in significantly faster communication between nodes.

3) Actuation: The order of AHU actuation was modified in Test 2 to allow for device settling time and prevent interference. In particular, in Tests 0 and 1, individual AHUs were ordered and actuated using a protocol that was not cognizant of settling times or building groupings, while the protocol was revised in Test 2 to systematically command the entire population of AHUs in a manner which maximized time between consecutive actuations for an individual unit.

Test 2 also featured a two-stage approach of actuation that was a result of the DER tracking behavior in Test 1. Some DERs, such as BESS, V1G EVs and V2G EVs, tracked quickly and accurately, whereas others, such as AHUs, tracked poorly. The overall tracking performance in Test 2 was improved by using "well-behaved" DERs to compensate for AHU tracking errors by incorporating the error signal from actuating AHUs in Stage 1 to the cumulative target signal for BESS, V1G EVs and V2G EVs in Stage 2. Although synchronous actuation of all participating DERs is preferred in practice, the two-stage approach highlights the significance of systematic characterization of DERs in minimizing ACE.

V. Test Results

A. Distributed Optimization/Cyber-Layer Results

In Table III, we present MSE results of our 1 s real-time Raspberry-pi distributed optimization solutions (the "cyber-layer" of the system).

TABLE III: Normalized mean-squared-error of distributed solutions obtained from real-time 1 s intervals compared to centralized solver solution for Test 2 (Section III-G)

DER Type	RC	PD	DANA	all
AHU	0	1.4×10^{-7}	2.8×10^{-9}	4.6×10^{-8}
V1G EVs	0	7.0×10^{-8}	1.7×10^{-9}	2.3×10^{-8}
V2G EVs	0	6.6×10^{-5}	5.0×10^{-7}	2.1×10^{-5}
BESS	0	2.0×10^{-6}	9.1×10^{-8}	6.5×10^{-7}
Total	0	1.8×10^{-5}	1.1×10^{-7}	4.9×10^{-6}

RC converged to the exact solution in all instances. This is unsurprising, as the RC problem formulation does not account for individual DER costs and thus, is a much simpler problem with a closed-form solution. For PD and DANA, we obtained excellent convergence, with errors on the order of 0.001% in the worst cases. In general, DANA tended to converge faster than PD in the sense that the obtained solutions were more accurate under the same fixed 1 s computation time. For our application with 1 s real-time windows, accuracy and convergence differences did not affect the physical layer results in any tangible way, but applications with more stringent accuracy or speed requirements may benefit from using a faster algorithm like DANA. The differences between DER populations can be largely attributed to the faster time scale of the V2G EVs (and to a lesser extent the BESS), see Table II. Since the V2G EVs were responsible for the high-frequency component of P_{ref}, the solver was required to converge to new solutions at every time step, which induced more error compared to the slow V1G EVs and AHUs with relatively static solutions.

B. Physical-Layer Test Results

We now present the results of the tracking performance pertaining to the physical-layer of the experiment. We provide only some selective plots for Test 0 and Test 1 in Fig. 3, and a complete set of plots for each Test 2 DER population in Fig. 4. Error and tracking delay data defined in Section III-G is given in Table IV for Test 1 and Test 2. Data for Test 0 is omitted due to its preliminary nature. The optimal shift described in Section III-G is applied to each time series and hence some areas in plots may appear like the provided signal anticipated the target.

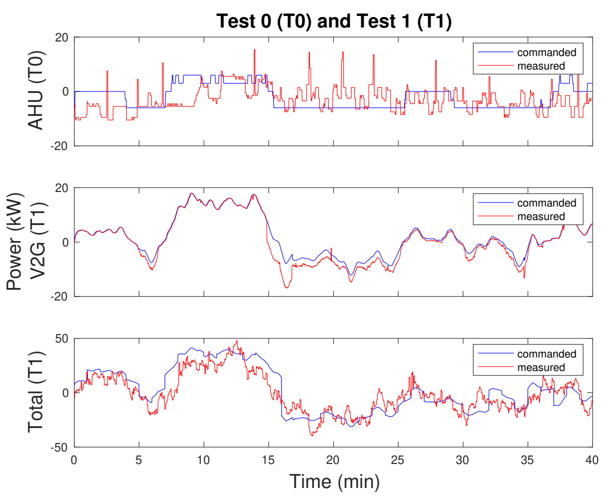

Fig. 3: Selected plots from Tests 0 and 1. **Top:** AHU response in Test 0. Note the poor tracking and spikes in the measured response. **Middle:** V2G response in Test 1. Note the inaccuracy in tracking during discharge-to-grid phases. **Bottom:** Total response in Test 1. Note the large-magnitude, low-frequency features demonstrating some broad tracking behavior, but overall poor performance.

Signal tracking accuracy in Test 0 was generally poor despite the small number of DERs employed, largely due to inexperience in actuating the AHUs and V1Gs. In particular, Fig. 3 reveals some oscillations in the AHU response. It is overall difficult to determine if even large-feature, low-frequency components of the signal were tracked. Further, data gathering for V1Gs and AHUs was done via noisy and unreliable building ION meters, which motivated the need for outlier treatment (Section III-F) in Tests 1 and 2, and resulted in the smoother and better tracking signal in the top plot of Fig. 4.

Test 1 yielded a 111% RMSE for AHUs. We speculate that the small 4 s delay in Test 1 is not representative of the actual AHU delay due to random correlations dominating the time shift for this large error. This is confirmed by a much better AHU response in Test 2 with RMSE 12%, where a 105 s

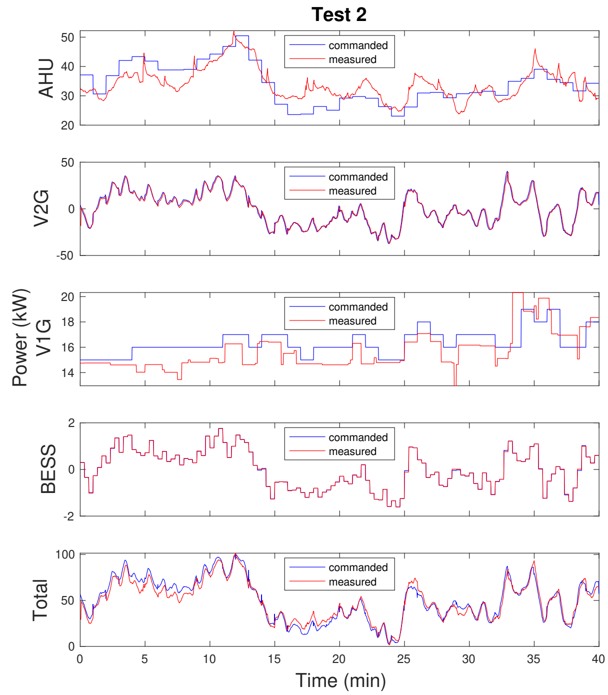

Fig. 4: Test 2 results. From **top to bottom**, AHU, V2G EVs, V1G EVs, BESS, and total responses. Note the substantially improved AHU, V2G, and total tracking performance compared to Figure 3.

delay is more likely to be representative of the true AHU actuation delay. Given the poor visibility into AHU and V1G controllers explained in Section IV, it is challenging to identify the source of the poor tracking behavior. We speculate that DER metering at the building level rather than the DER level was a major source of error for AHU and V1G in Test 1. This was largely resolved in Test 2 by utilizing a different population of V1Gs with dedicated meters and by modifying the actuation scheme for AHUs to be less susceptible to metering errors as described in Section IV-D. Additionally, the actuation-interface stalling for V1G EVs, described in Section IV-C, was dominant in Test 1, resulting in the poor tracking for V1Gs. Actuating-interface issues were resolved in Test 2 by utilizing an automated control scheme for the V1Gs, which led to significantly lower error.

The BESS emerged as the star performer achieving very accurate tracking across all tests with no delay. The V2G EVs also performed relatively well aside from a signal overshoot issue observed during the discharge cycle in Test 1 seen in Fig. 3. The issue was resolved in Test 2 by using V2G EV charging stations from a different manufacturer (Princeton Power), as described in Section IV-D. The V2G charging stations deployed for these tests were pre-commercial or early commercial models that had a few operating issues, such as the overshoot issue during Test 1.

The inability of the AHUs to respond to steep, short ramps (Fig. 4) could be due to slow start-up sequences programmed into the building automation controllers to increase device longevity or due to transients associated with driving their AC induction electric motors. Tackling this would require dynamic models and parameter identification of signal response and delay. With the new V1G EV population in Test 2, tracking delay reduced from 40 s to 10 s and the tracking accuracy improved significantly. The 1 kW bias seen in Fig. 4 is likely due to rounding errors arising from the inability of PowerFlex charging stations to accept non-integer setpoints.

The superior performance of the BESS and V2Gs motivated the two-stage actuation scheme described in Section IV-D, which contributed to reducing the total RMSE from 50% in Test 1 to 10% in Test 2 (compare the bottom plots of Figs 3 and 4). The two-stage approach allows a sufficiently large proportion of accurately tracking DERs to compensate for the errors of the first stage, where tracking is worse. In this way, poorly-tracking DERs, such as AHUs, can still contribute by loosely tracking some large-feature, low-frequency components of the target signal. The low-frequency contribution reduces the required total capacity of the strongly-performing DERs in the second stage leading to more fine-tuned signal tracking in aggregation. Some recommended rules of thumb for two-stage approach are: (i) Total capacity of first-stage DERs is less than or equal to total capacity of second-stage DERs. (ii) DERs in the first stage are capable of tracking with $<$ 50% RMSE. (iii) DER cost functions are such that the deviation from the baseline is lower cost for first-stage DERs than for second-stage. (iii) allocates a significant portion of the target signal initially to first-stage DERs, freeing up DER capacity in the second-stage for error compensation.

TABLE IV: **Left:** Relative root mean-squared-error of tracking error by DER type. **Right:** Delay (optimal time-shift) of DER responses in seconds.

DER Type	Test 1	Test 2
AHU	1.11	0.12
V1G EVs	0.68	0.077
V2G EVs	0.30	0.060
BESS	0.054	0.018
Total	0.50	0.097

DER Type	Test 1	Test 2
AHU	4	105
V1G EVs	40	10
V2G EVs	5	3
BESS	0	0
Total	N/A	N/A

C. Economic Benefit Analysis

Here, we evaluate the economic benefit of the proposed test system, which is vital for wider scale adoption of DERs as a frequency regulation resource in real electricity markets. To this end, we take an approach similar to [10] to first demonstrate that the testbed is eligible to participate in the PJM ancillary services market. Following the PJM Manual 12 [28] (Section III-G), we compute a Correlation Score S_c = 0.98, Delay Score S_d = 0.65, and Precision Score S_p = 0.91 from data for Test 2, and obtain a Performance Score $S = 0.85 \geq 0.75$, which confirms the eligibility to participate in the PJM ancillary service market.

Next, we compute the estimated annual revenue assuming that the resources are available throughout the day. Using PJM's ancillary service market data[2] with our total (active) DER capacity of 184 kW and performance score of 0.85, the capability and performance credits for this population of resources (cf. [29, Section 4]) would respectively be \$135 and \$11, for July 9, 2020. This gives an estimated amount of \$53,290 as the total annual revenue. Note that the 184 kW DER capacity employed in this work represents less than 5% of the total DER capacity and less than 0.5% of the total capacity of the UCSD microgrid, cf. [18]. As such, the revenue would significantly increase if more microgrid resources are utilized for regulation, even with reduced availability.

VI. Conclusions

We have presented one of the first real-world demonstrations of secondary frequency response in a distribution grid using up to 176 heterogeneous DERs. The DERs include AHUs, V1G and V2G EVs, a BESS, and passive building loads and PV generators. The computation setup utilizes state-of-the-art distributed algorithms to find the solution of a power allocation problem. We show that the real-time distributed solutions are close to the true centralized solution in an MSE sense. Tests with real, controllable DERs at power closely track the given active-power reference signal in aggregation. Further, our economic benefit analysis shows a potential annual revenue of \$53K for the chosen DER population. These tests highlight the importance of dedicated and noise-free measurement sensors and a well-understood and reliable DER control interface for precise signal tracking. Extensions of this work are ongoing under DERConnect[3], a new project at UCSD that aims to develop a testbed consisting of 2500 DERs that allows for online implementation of various distributed algorithms. As is already recognized by the power systems community and federal funding agencies such as ARPA-e and National Science Foundation, large-scale power-in-the-loop testing is needed for transitioning distributed technologies to real distribution systems. We hope that this work spurs further testing and ultimately widespread adoption of coordinated resource control algorithms by relevant players in industry.

Acknowledgements

We would like to thank numerous people in the UCSD community and beyond for their generous contributions of time and resources to enable such an ambitious project to come together. We extend thanks to: (i) Aaron Ma and Jia (Jimmy) Qiu for assisting with hardware setup and software development for the distributed computation systems; (ii) Kevin Norris for coordinating the fleet vehicles; (iii) Abdulkarim Alamad for overseeing V1G drivers in Test 2; (iv) Kelsey Johnson for managing the Nuvve contributions; (v) Ted Lee, Patrick Kelly, and Steven Low for managing the PowerFlex contribution; (vi) Marco Arciniega, Martin Greenawalt, James Gunn, Josh Kavanagh, Jennifer Rodgers, Patricia Roman and Lashon Smith from UCSD parking for reserving EV charging station parking spaces; (vii) Charles Bryant, Harley Crace, John Denhart, Nirav Desai, John Dilliott, Mark Gaus, Martin Greenawalt, Gerald Hernandez, Brandon Hirsch, Mark Jurgens, Josh Kavanagh, Jose Moret, Chuck Morgan, Curt Lutz, Jose Moret, Cynthia Wade, Raymond Wampler and Ed Webb for contributing their EVs in Test 1; (viii) Adrian Armenta, Adrian Gutierrez and Minghua Ong who helped with ChargePoint manual control; (ix) Bob Caldwell (Centaurus

[2]https://dataminer2.pjm.com/feed/reg_prices/definition

[3]https://sites.google.com/ucsd.edu/derconnect/home

Prime), Gregory Collins, Charles Bryant, and Robert Austin for programming and enabling the AHU control; (x) Gary Matthews and John Dilliott for permitting the experimentation on "live" buildings and vehicles; and (xi) Antoni Tong and Cristian Cortes-Aguirre for supplying the BESS. Finally, we would like to extend a sincere thanks to the ARPA-e NODES program for its financial support and to its leadership, including Sonja Glavaski, Mario Garcia-Sanz, and Mirjana Marden, for their vision and push for the development of large-scale power-in-the-loop testing environments.

REFERENCES

[1] Z. A. Obaid, L. M. Cipcigan, L. Abrahim, and M. T. Muhssin, "Frequency control of future power systems: reviewing and evaluating challenges and new control methods," *Journal of Modern Power Systems and Clean Energy*, vol. 7, no. 1, pp. 9–25, 2019.

[2] M. Kintner-Meyer, "Regulatory policy and markets for energy storage in North America," *Proceedings of the IEEE*, vol. 102, no. 7, pp. 1065–1072, 2014.

[3] CAISO, "Pay for performance regulation: Draft final proposal addendum," February 22 2012. [Online]. Available: http://www.caiso.com/Documents/Addendum-DraftFinalProposal-Pay_PerformanceRegulation.pdf

[4] P. Srivastava, C.-Y. Chang, and J. Cortés, "Participation of microgrids in frequency regulation markets," in *American Control Conference*, Milwaukee, WI, May 2018, pp. 3834–3839.

[5] B. Kroposki, "Basic research needs for autonomous energy grids-Summary report of the workshop on autonomous energy grids," NREL, Tech. Rep., September 13-14 2017, NREL/TP-5D00-70428.

[6] W. Kempton, V. Udo, K. Huber, K. Komara, S. Letendre, S. Baker, D. Brunner, and N. Pearre, "A test of vehicle-to-grid (V2G) for energy storage and frequency regulation in the PJM system," 2008, available at http://www1.udel.edu/V2G/resources/test-v2g-in-pjm-jan09.pdf.

[7] M. Swierczynski, D. Stroe, A. Stan, R. Teodorescu, R. Lrke, and P. C. Kjr, "Field tests experience from 1.6MW/400kWh Li-ion battery energy storage system providing primary frequency regulation service," in *IEEE PES ISGT Europe*, 2013, pp. 1–5.

[8] M. Cheng, J. Wu, S. J. Galsworthy, C. E. Ugalde-Loo, N. Gargov, W. W. Hung, and N. Jenkins, "Power system frequency response from the control of bitumen tanks," *IEEE Transactions on Power Systems*, vol. 31, no. 3, pp. 1769–1778, 2016.

[9] J. Liu, W. Zhang, and Y. Liu, "Primary frequency response from the control of led lighting loads in commercial buildings," *IEEE Transactions on Smart Grid*, vol. 8, no. 6, pp. 2880–2889, 2017.

[10] Y. Lin, P. Barooah, S. Meyn, and T. Middelkoop, "Experimental evaluation of frequency regulation from commercial building HVAC systems," *IEEE Transactions on Smart Grid*, vol. 6, no. 2, pp. 776–783, 2015.

[11] E. Vrettos, E. C. Kara, J. MacDonald, G. Andersson, and D. S. Callaway, "Experimental demonstration of frequency regulation by commercial buildingspart II: Results and performance evaluation," *IEEE Transactions on Smart Grid*, vol. 9, no. 4, pp. 3224–3234, 2018.

[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 *2018 IEEE Power Energy Society General Meeting*, Portland, OR, Aug. 2018, pp. 1–5.

[13] L. Fabietti, T. T. Gorecki, F. A. Qureshi, A. Bitlislioğlu, I. Lymperopoulos, and C. N. Jones, "Experimental implementation of frequency regulation services using commercial buildings," *IEEE Transactions on Smart Grid*, vol. 9, no. 3, pp. 1657–1666, 2018.

[14] K. Baker, X. Jin, , D. Vaidhynathan, W. Jones, D. Christensen, B. Sparn, J. Woods, H. Sorensen, and M. Lunacek, "Frequency regulation services from connected residential devices: Short paper," in *Proceedings of the 3rd ACM International Conference on Systems for Energy-Efficient Built Environments*, Palo Alto, CA, 2016, pp. 119–122.

[15] SDGE, "AC saver for business." [Online]. Available: https://www.sdge.com/businesses/savings-center/energy-management-programs/demand-response/summer-saver-program

[16] S. Cherian and P. Asmus, "Liberating microgrids (and all DER): Aligning customer needs with solutions provider offerings," White Paper, Navigant Research, 2016.

[17] A. Tuckey, S. Zabihi, and S. Round, "Decentralized control of a microgrid," in *European Conference on Power Electronics and Applications*, Warsaw, Poland, Sept. 2017, pp. 1–10.

[18] B. Washom, J. Dilliot, D. Weil, J. Kleissl, N. Balac, W. Torre, and C. Richter, "Ivory tower of power: Microgrid implementation at the University of California, San Diego," *IEEE Power and Energy Magazine*, vol. 11, no. 4, pp. 28–32, 2013.

[19] "PJM markets and operations: Anciliary services," https://www.pjm.com/markets-and-operations/ancillary-services.aspx, 2019, regD Normalized Signal Test (after 1.30.2017).

[20] A. D. Domínguez-García, C. N. Hadjicostis, and N. H. Vaidya, "Resilient networked control of distributed energy resources," *IEEE Journal on Selected Areas in Communications*, vol. 30, no. 6, pp. 1137–1148, July 2012.

[21] A. Cherukuri and J. Cortés, "Distributed algorithms for convex network optimization under non-sparse equality constraints," in *Allerton Conf. on Communications, Control and Computing*, Monticello, IL, Sept. 2016, pp. 452–459.

[22] A. Cherukuri, B. Gharesifard, and J. Cortés, "Saddle-point dynamics: conditions for asymptotic stability of saddle points," *SIAM Journal on Control and Optimization*, vol. 55, no. 1, pp. 486–511, 2017.

[23] T. Anderson, C.-Y. Chang, and S. Martínez, "Distributed approximate Newton algorithms for constrained optimization," *Automatica*, vol. 109, p. 108538, 2019.

[24] T. Anderson, M. Muralidharan, P. Srivastava, H. V. Haghi, J. Cortés, J. Kleissl, S. Martínez, and B. Washom, "Frequency regulation with heterogeneous energy resources: A realization using distributed control," *arXiv preprint arXiv:2007.07971*, 2020.

[25] CAISO, "Vehicle-grid integration (vgi) roadmap: Enabling vehicle-based grid services," February 2014. [Online]. Available: http://www.caiso.com/documents/vehicle-gridintegrationroadmap.pdf

[26] ——, "Business practice manual for market operation," May 24 2018, version 57. [Online]. Available: https://bpmcm.caiso.com/Pages/BPMDetails.aspx?BPM=Market\%20Operations

[27] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 2.1," http://cvxr.com/cvx, Mar. 2014.

[28] "PJM manual 12: Balancing operations," March 26 2020, revision 40. [Online]. Available: https://www.pjm.com/-/media/documents/manuals/m12.ashx

[29] "PJM manual 28: Operating agreement accounting," December 3 2019, revision 83. [Online]. Available: https://www.pjm.com/-/media/documents/manuals/m28.ashx