Real Time-Optimal Power Flow-Based Distributed Energy Resource Management System (DERMS)

Jing Wang, Joshua Comden, and Andrey Bernstein

National Renewable Energy Laboratory

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Technical Report
NREL/TP-5D00-87767
May 2024

Real Time-Optimal Power Flow-Based Distributed Energy Resource Management System (DERMS)

Jing Wang, Joshua Comden, and Andrey Bernstein

National Renewable Energy Laboratory

Suggested Citation

Wang, Jing, Joshua Comden, and Andrey Bernstein. 2024. *Real Time-Optimal Power Flow Based Distributed Energy Resource Management System (DERMS)*. Golden, CO: National Renewable Energy Laboratory. NREL/TP-5D00-87767. https://www.nrel.gov/docs/fy24osti/87767.pdf.

Contract No. DE-AC36-08GO28308

Technical Report
NREL/TP-5D00-87767
May 2024

National Renewable Energy Laboratory
15013 Denver West Parkway
Golden, CO 80401
303-275-3000 • www.nrel.gov

NOTICE

This work was authored 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 U.S. Department of Energy Office of Energy Efficiency and Renewable Energy Solar Energy Technologies Office. The views expressed herein do not necessarily represent the views of the DOE or the U.S. Government.

Cover Photos by Dennis Schroeder: (clockwise, left to right) NREL 51934, NREL 45897, NREL 42160, NREL 45891, NREL 48097, NREL 46526.

Acknowledgments

The authors thank Georgios Stefopoulos of the U.S. Department of Energy's Solar Energy Technologies Office for his continuous support and very helpful insights for this project.

In addition, we acknowledge the substantial contributions to this project provided by staff at NVIDIA and Utilidata Inc., including Steven Forsyth and Reynaldo Gomez from NVIDIA, and Ning Li, Taylor Spalt, and Marissa Hummon from Utilidata Inc.

Finally, the authors acknowledge the substantial contribution to this project provided by Subhankar Ganguly, who supported the hardware-in-the-loop testing.

List of Acronyms

ADMS	advanced distribution management system
DER	distributed energy resource
DERMS	distributed energy resource management system
DNP3	Distributed Network Protocol 3
EVSE	electric vehicle supply equipment
HELICS	Hierarchical Engine for Large-Scale Infrastructure Co-Simulation
HIL	hardware-in-the-loop
LTC	load tap changer
NREL	National Renewable Energy Laboratory
PCC	point of common coupling
PHIL	power-hardware-in-the-loop
PV	photovoltaic
RT-OPF	real-time optimal power flow
VPP	virtual power plant

Executive Summary

The smart meter-based real-time optimal power flow (RT-OPF) distributed energy resource management system (DERMS) is a technology that monitors, controls, and coordinates large numbers of distributed energy resources (DERs) in real time to provide aggerated grid services to the electric utility and to integrate customers' preferences. Although many DERMS solutions have been developed recently, they typically operate on a slow time scale; use centralized, computationally heavy optimization methods; and disregard power system operational constraints. In contrast, the National Renewable Energy Laboratory's (NREL's) meter-based RT-OPF DERMS provides a unique real-time, distributed, and plug-and-play optimization platform that coordinates the operation of massive numbers of DERs to ensure voltage and power quality, to maximize social welfare, and to emulate virtual power plants (VPPs). The distributed nature of the platform allows scalability to millions of DERs using low-power microcontrollers running applications locally (e.g., in residential homes, substations, and service transformers). These salient features create mature conditions for RT-OPF DERMS to be widely adopted and commercialized. The commercialization path focuses on embedding the technology in utility infrastructure via the electricity meter. This pathway capitalizes on utility hardware and communications investments and accelerates progress toward grid modernization goals, utilities' needs for control of DERs, and end-user customer benefits.

This technology has the potential to *double* the hosting capacity of distribution systems, reduce the system upgrade costs by 10% or more, yield economic value and potential savings by 5%, enable the growth of the clean energy economy, and reduce fossil-fuel carbon emissions.

Therefore, the objective of this project is to promote the scalability and commercialization of RT-OPF DERMS. Thus far, this technology has shown promising results when applied in multiple demonstrations in the lab and field. Based on our experiences in multiple projects, we propose three technical aspects to be improved in this project for better commercialization. Therefore, the metrics of success are: (1) develop automated tuning and deployment processes with guaranteed performance and stability; (2) investigate the limits of RT-OPF under various communication scenarios and analyze the robustness of RT-OPF in terms of communication failures; and (3) demonstrate that RT-OPF works in the control platform that can be embedded in next-generation smart meters. Note that utilities are commercially motivated to invest in infrastructure, and regulators are trying to ensure that utilities make prudent investments that prepare them for a dynamic, DER-heavy grid. If we can technically integrate software that prepares utilities for future integration into their existing hardware that the utilities are already proposing to deploy at scale, we will have developed an incredibly scalable pathway to market.

This technology represents a paradigm shift in DER integration. RT-OPF installed on the smart meter establishes a grid-edge platform capable of planning, integrating, and optimizing DERs. This hardware-embedded path to market will have a much higher chance of driving results and associated economic growth than simply selling a software product to utilities and/or attempting to deploy special-purpose hardware to enable the solution. More importantly, the RT-OPF technology itself outperforms existing solutions by being decentralized and real-time, using low-power computation, employing standard implementation to reduce deployment costs, and accurately controlling performance so all entities experience increased performance. If proven successful, this technology will be game-changing and could result in millions of dollars in

additional cost reduction for utilities and hundreds of dollars savings in electricity bills for customers per year.

Table of Contents

List of Figures

List of Tables

1 Project Overview

1.1 Background

With the proliferation of distributed energy resources (DERs) in distribution systems, many utilities face challenges with integrating large numbers of non-utility devices into operations at all levels and ensuring grid stability and controls over them. Industry and Institute of Electrical and Electronics Engineers working groups have begun to address DER integration by identifying, standardizing, and requiring the functions that individual DERs can perform autonomously, such as frequency-watt and voltage-volt ampere reactive (volt-var) controls (Wang et al. 2020b). However, in many situations, significantly improved performance can be achieved with feeder-wide coordination systems such as distributed energy resource management systems (DERMS), as shown in Figure 1. DERMS have been an emerging technology that seamlessly integrate DERs into distribution grids at the feeder level in a coordinated manner, and they provide multiple benefits to both utilities and customers in a win-win situation, bridging the gap between DER group-managing entities (e.g., DMS) and DER devices to aggregate (Seal, Renjit, and Deaver 2018).

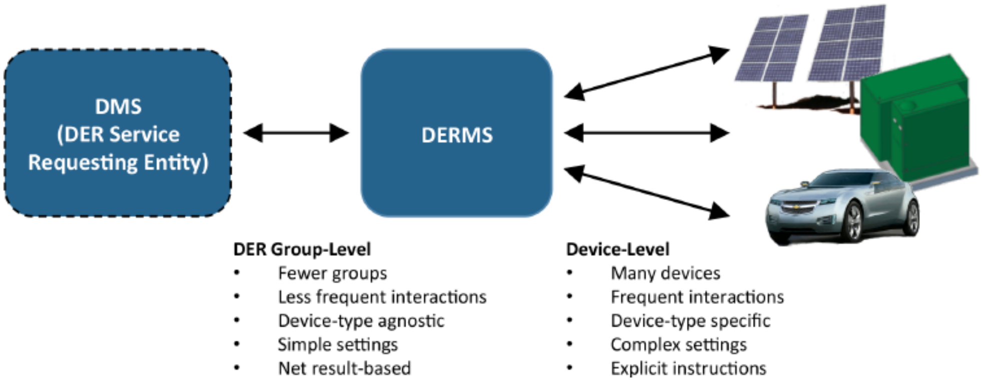

Figure 1. DERMS functionalities

(Seal, Renjit, and Deaver 2018)

Although many DERMS platforms have been developed recently, they typically operate on a slow time scale; use centralized, computationally heavy optimization methods; and disregard power system operational constraints. In contrast, the National Renewable Energy Laboratory's (NREL's) real-time optimal power flow (RT-OPF) DERMS provides a unique *real-time*, *distributed*, and *plug-and-play* optimization platform that coordinates the operation of massive numbers of DERs to ensure *voltage and power quality*, *maximize social welfare*, and *emulate virtual power plants* (VPPs) (Bernstein et al. 2020). The distributed nature of the platform allows scalability to millions of DERs using low-power microcontrollers running standard code locally (e.g., in residential homes, substations, and service transformers). These salient features create mature conditions for RT-OPF DERMS to be widely adopted and commercialized.

1.2 Project Objectives

This technology commercialization project aims to improve the maturity of a lab-proven clean energy technology to a commercial ready technology for product development, integrate the technology with broader systems, provide extended performance data, and validate the manufacturability and reliability of the technology. The lab-proven technology, RT-OPF DERMS, was developed and validated through previous U.S. Department of Energy-funded efforts, including Advanced Research Projects Agency–Energy funding under the Network Optimized Distributed Energy Systems program and Holy-Cross Energy High Impact Project. In the Advanced Research Projects Agency–Energy Network Optimized Distributed Energy Systems project, the RT-OPF DERMS was developed and implemented in multiple hardware platforms, demonstrating its performance and capabilities in the lab and field environments. The technology was also evaluated and matured via a participation in the U.S. Department of Energy I-Corps program, whose goal is to pair teams of researchers with industry mentors for an intensive 2-month training in which the researchers define technology value propositions, conduct customer discovery interviews, and develop viable market pathways for their technologies. These activities indicate the high technology maturity and Technology Readiness Level of the RT-OPF DERMS.

However, the single-use hardware leveraged in the previous implementations limits the scalability and commercial potential of the RT-OPF DERMS. Therefore, NREL collaborates with Utilidata Inc. and NVIDIA to integrate the technology into a controller platform (the DER chip provided by NVIDIA) that can be embedded in the next-generation smart meters, which will be installed by many utilities in the near future. Utilidata has already purchased an option agreement for the patent associated with this technology and is evaluating the possibility of licensing. Success in this project would hatch a next-generation smart meter with control intelligence for DER management, thus enabling industry, investors, and partners to justify substantial commitments of financial resources, personnel, production facilities, and materials to develop promising technologies into early commercial products.

1.3 Project Summary

The technology commercialization fund project represents a joint effort between NREL, Utilidata Inc., and NVIDIA to improve the RT-OPF DERMS technology maturity for wider deployment. We first used NREL's Integrated Grid Modeling System-Hierarchical Engine for Large-Scale Infrastructure Co-Simulation (HELICS) to test the automatic tuning and perform a communication impact study, then we performed laboratory controller and power-hardware-in-the-loop (PHIL) testing. This project achieves the following outcomes when it is completed:

- Improve the RT-OPF DERMS technology maturity by developing the automatic tuning of the RT-OPF algorithm and testing under various scenarios and dynamic grid operating conditions
- Investigate the RT-OPF performance (voltage regulation and VPP) under various communication issues and identify the requirements for communication for deploying RT-OPF in the field
- Define the RT-OPF communication and control architecture for real-world applications based on the inputs from the advisory board (inverter vendors)
- Integrate the RT-OPF DERMS local control algorithm in the NVIDIA DER chip and test the DER chip's communication and control performance in a controller and PHIL environment to mimic the field deployment.

2 Technology Maturity Development

To make the RT-OPF DERMS ready for commercialization and product development, we will improve the technology maturity from three aspects: automatic tuning of RT-OPF DERMS, study the communication issues and define the communication requirements, and embed the RT-OPF algorithm in the DER chip and evaluate through hardware-in-the-loop (HIL) with real hardware inverters. The first two are essentially needed to improve the technology maturity for commercialization, and the third aspect is to directly evaluate the performance of the DER chip which will be inserted into next-generation smart meter for DER management. The development of each aspect will be explained in detail.

2.1 Automatic Tuning of RT-OPF DERMS

To describe the theory behind the automatic-tuning method designed and tested in this project, the system model and RT-OPF DERMS algorithm from (Bernstein and Dall'Anese 2017) was borrowed and modified as shown in Figure 2. We first briefly describe the RT-OPF DERMS algorithm and the modifications made to allow for an auto-tuning method to be added on. Then we describe the auto-tuning method we designed and tested.

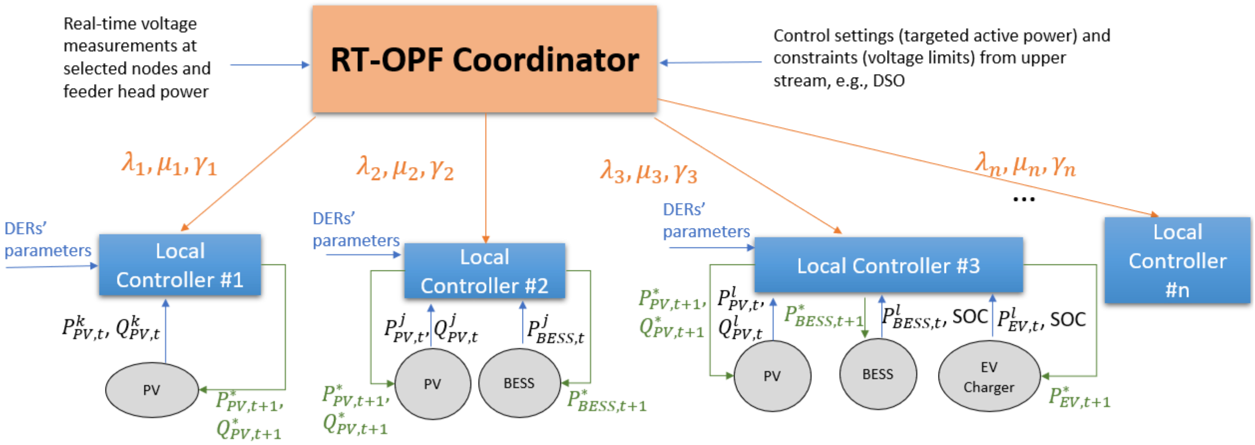

Figure 2. Control architecture of the RT-OPF for DERMS application

2.1.1 RT-OPF DERMS

The general flow of information for the RT-OPF DERMS algorithm is that the coordinator (see Figure 3) receives:

- $\hat{\boldsymbol{v}}$ measurements of the voltage magnitudes from the set of measurement buses
- $(\underline{\boldsymbol{P}}_0, \overline{\boldsymbol{P}}_0)$ the VPP active power band set points for each phase
- $\widehat{\boldsymbol{P}}_0$ active power measurement from each phase at the feeder head

calculates and sends out a set of gradients to each distributed controller of a DER on the feeder. Each distributed controller (see Figure 4) receives:

- A set of gradients from the coordinator (λ, μ and γ)
- A measured active and reactive power injection of the DER is tasked to control

- Other measurements, such as available photovoltaic (PV) active power and battery state of charge.

calculates and sends back to the DER an active and reactive power set point. The DER implements its set point on the feeder, which collectively regulates the voltage magnitudes and feeder head power within the target limits/boundaries.

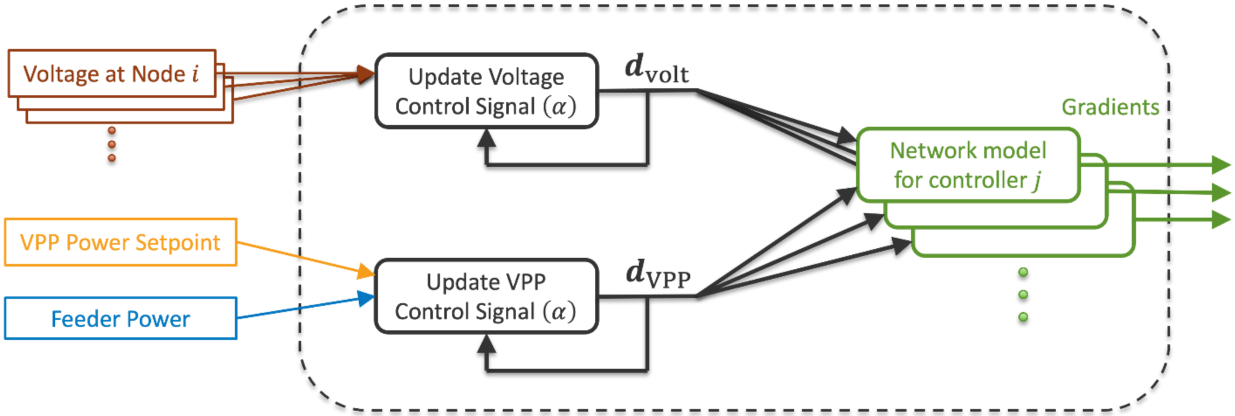

Figure 3. Input/output diagram of the RT-OPF DERMS coordinator

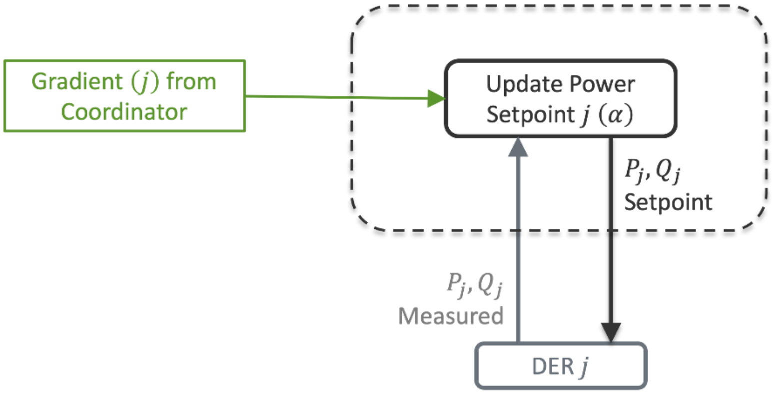

Figure 4. Input/output diagram of an RT-OPF DERMS distributed coordinator for DER j

The dual variables on voltage regulation $\boldsymbol{d}_{volt}(t) \coloneqq \left[\underline{\boldsymbol{d}}_{volt}(t)^{\intercal} \quad \overline{\boldsymbol{d}}_{volt}(t)^{\intercal}\right]^{\intercal}$ are partitioned into two parts: $\underline{\boldsymbol{d}}_{volt}(t)$ are the dual variables for violating the lower bounds on the voltage magnitudes $\underline{\boldsymbol{v}}$; and $\overline{\boldsymbol{d}}_{volt}(t)$ are the dual variables for violating the upper bounds on the voltage magnitudes $\overline{\boldsymbol{v}}$. The update procedure on both sets of variables performed by the coordinator takes the measured voltages $\widehat{\boldsymbol{v}}(t-1)$ as an input with the following equations:

$$\underline{\boldsymbol{d}}_{volt}(t) \coloneqq proj_{\mathbb{R}_+}\left\{\underline{\boldsymbol{d}}_{volt}(t-1) + \alpha\left(\underline{\boldsymbol{v}} - \widehat{\boldsymbol{v}}(t-1) - \epsilon\underline{\boldsymbol{d}}_{volt}(t-1)\right)\right\} \tag{1a}$$

$$\overline{\boldsymbol{d}}_{volt}(t) \coloneqq proj_{\mathbb{R}_+}\left\{\overline{\boldsymbol{d}}_{volt}(t-1) + \alpha\left(\widehat{\boldsymbol{v}}(t-1) - \overline{\boldsymbol{v}} - \epsilon\overline{\boldsymbol{d}}_{volt}(t-1)\right)\right\} \tag{1b}$$

where $proj_{\mathbb{R}_+}$ is the projection onto the non-negative real number space, and α and ϵ are positive parameters. Notice that an increase in any of the elements of $\underline{\boldsymbol{d}}_{volt}(t)$ or $\overline{\boldsymbol{d}}_{volt}(t)$ comes from a measured violation of either voltage lower or upper bound, respectively. A decrease comes from the term with ϵ, which is meant to proportionally lower the value of the elements of $\underline{\boldsymbol{d}}_{volt}(t)$ and $\overline{\boldsymbol{d}}_{volt}(t)$ toward zero when voltage violations are not occurring.

Similarly, the dual variables for the VPP $\boldsymbol{d}_{VPP}(t) \coloneqq \left[\underline{\boldsymbol{d}}_{VPP}(t)^{\intercal} \quad \overline{\boldsymbol{d}}_{VPP}(t)^{\intercal}\right]^{\intercal}$ are portioned into two parts: $\underline{\boldsymbol{d}}_{VPP}(t)$ are the dual variables for the time-varying lower bound on the VPP band $\underline{\boldsymbol{P}}_0(t)$; and $\overline{\boldsymbol{d}}_{VPP}(t)$ are the dual variables for the time-varying upper bound on the VPP band $\overline{\boldsymbol{P}}_0(t)$. The update procedure on both sets of variables performed by the coordinator takes the measured active feeder power of each phase $\widehat{\boldsymbol{P}}_0(t-1)$ as an input with the following equations:

$$\underline{\boldsymbol{d}}_{VPP}(t) \coloneqq proj_{\mathbb{R}_+}\left\{\underline{\boldsymbol{d}}_{VPP}(t-1) + \alpha\left(\underline{\boldsymbol{P}}_0(t) - \widehat{\boldsymbol{P}}_0(t-1) - \epsilon\underline{\boldsymbol{d}}_{VPP}(t-1)\right)\right\} \tag{2a}$$

$$\overline{\boldsymbol{d}}_{VPP}(t) \coloneqq proj_{\mathbb{R}_+}\left\{\overline{\boldsymbol{d}}_{VPP}(t-1) + \alpha\left(\widehat{\boldsymbol{P}}_0(t-1) - \overline{\boldsymbol{P}}_0(t) - \epsilon\overline{\boldsymbol{d}}_{VPP}(t-1)\right)\right\}. \tag{2b}$$

To transform the dual variables into individualized gradients for each controller, we employ a linearization of the power flow equations, which gives the following approximated linear relationship between the active and reactive power injections $(\boldsymbol{p}, \boldsymbol{q})$ and the voltage magnitudes $\boldsymbol{v}$ and the feeder active powers $\boldsymbol{P}_0$:

$$\boldsymbol{A}\boldsymbol{p} + \boldsymbol{B}\boldsymbol{q} + \boldsymbol{a} \approx \boldsymbol{v} \tag{3a}$$

$$\boldsymbol{M}\boldsymbol{p} + \boldsymbol{N}\boldsymbol{q} + \boldsymbol{o} \approx \boldsymbol{P}_0. \tag{3b}$$

The matrices and vectors of linearization coefficients $(\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{a}, \boldsymbol{M}, \boldsymbol{N}, \boldsymbol{o})$ can be derived from the nodal admittance matrix (see Bernstein et al. [2020] as an example). With matrices of linearization coefficients $(\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{M}, \boldsymbol{N})$, the dual variables $\left(\boldsymbol{d}_{volt}(t), \boldsymbol{d}_{VPP}(t)\right)$ can be transformed into the following gradients:

$$\boldsymbol{g}_P(t) \coloneqq \boldsymbol{A}^{\intercal}\left(\overline{\boldsymbol{d}}_{volt}(t) - \underline{\boldsymbol{d}}_{volt}(t)\right) + \boldsymbol{M}^{\intercal}\left(\overline{\boldsymbol{d}}_{VPP}(t) - \underline{\boldsymbol{d}}_{VPP}(t)\right) \tag{4a}$$

$$\boldsymbol{g}_Q(t) \coloneqq \boldsymbol{B}^{\intercal}\left(\overline{\boldsymbol{d}}_{volt}(t) - \underline{\boldsymbol{d}}_{volt}(t)\right) + \boldsymbol{N}^{\intercal}\left(\overline{\boldsymbol{d}}_{VPP}(t) - \underline{\boldsymbol{d}}_{VPP}(t)\right). \tag{4b}$$

The elements of the gradients give the negative direction that each bus should shift its power injection to get the voltage magnitudes and VPP within their respective bounds. The elements of the gradients are sent to the distributed controllers associated with their location in the feeder.

The distributed controller j receives its individualized gradient $\left(g_{P,j}(t), g_{Q,j}(t)\right)$ from the coordinator and power injection measurements $\left(\hat{P}_j(t-1), \hat{Q}_j(t-1)\right)$ from its associated DER j. The controller keeps track of a cost $c_j\left(\hat{P}_j(t-1), \hat{Q}_j(t-1); t-1\right)$ associated with its local objective (e.g., limit active power curtailment of a PV, keep a storage unit near a specified state of charge, etc.), and updates a power injection set point for DER j with the following equation:

$$\begin{bmatrix} P_j(t) \\ Q_j(t) \end{bmatrix} \coloneqq proj_{\mathcal{Y}_j(t-1)}\left\{\begin{bmatrix} \hat{P}_j(t-1) \\ \hat{Q}_j(t-1) \end{bmatrix} - \alpha\left(\nabla_{P,Q}c_j\left(\hat{P}_j(t-1), \hat{Q}_j(t-1); t-1\right) + \begin{bmatrix} g_{P,j}(t) \\ g_{Q,j}(t) \end{bmatrix} + \nu\begin{bmatrix} \hat{P}_j(t-1) \\ \hat{Q}_j(t-1) \end{bmatrix}\right)\right\} \tag{5}$$

where $\nabla_{P,Q}$ is the gradient operator with respect to the active and reactive power injections, $proj_{\mathcal{Y}_j(t-1)}$ is the projection onto the feasible active and reactive power injection space $\mathcal{Y}_j(t-1)$ of DER j, and ν is a positive parameter. This update moves the set point of DER j in a direction that either lowers its local cost, helps the DERMS coordinator meet the voltage regulation and VPP bounds, or both.

2.1.2 Automatic Tuning

One of the major time-intensive activities for an engineer to implement the system above is to tune the parameter α, which is the step size for updating the dual variables in Equations (1) and (2), and power injections (primal variables) in Equation (5). If the step size α is set too low, then RT-OPF DERMS reacts too slowly, which can cause the voltage and VPP bounds to not be satisfied. On the other hand, if the step size α is set too high, then the system can become unstable as the control actions become too sensitive to changes in the system measurements. Thus, the engineer must be conservative in choosing an effective value for the step size α that is far from the point that makes the system unstable. Another complicating factor is that a good value for the step size α depends on the conditions of the system, such as volatility in the PV generation, changes in the VPP band, and changes caused by a superseding advanced distribution management system (ADMS). Therefore, auto-tuning becomes the most critical technology improvement to make the NREL-developed RT-OPF DERMS able to be commercialized.

Based on extensive literature review and deep understanding of the fundamentals of the RT-OPF algorithm, our plan was to try the following approaches to do the automatic tuning, ordered from simple to more complex:

1. Acceleration (increase step size) and deceleration (decrease step size) based on the direction of the updates (Dall'Anese, Guggilam, and Simonetto 2017). The general idea is that if two consecutive updates (in either the primal or dual variables) are in the same direction, then the size of the updates should be increased to get to the point of convergence faster. Otherwise, if the updates are in different directions, then the size of the updates should be decreased to lessen or eliminate oscillations around the point of convergence.
2. Primal-dual balancing. (Goldstein, Li, and Yuan 2015) At each time step, either the primal or the dual variables may be relatively farther from their associated optimal convergence point. Increasing the weight of the step size associated with either the primal or dual variables that are further away from an estimated convergence point and decreasing the weight of the other set of variables will get to the optimal convergence point faster.
3. Spectral adaptive step size. (Obi, Slay, and Bass 2020) At each time step, fit a linear function to the gradient and find an optimal step size for the linearized system. If the true non-linear function deviates too much from the linearization, backtrack down the step size.

However, we found that the basic idea in (Dall'Anese, Guggilam, and Simonetto 2017) was sufficient for our need to auto-tune to the step size α. Dall'Anese, Guggilam, and Simonetto (2017) make an algorithm become adaptive that has a similar structure to RT-OPF DERMS for a completely different application. The idea is to increase or decrease the step size based on variations of the gradient directions, or in our case, the direction of the updates. If an update goes in the same direction as its previous update, it indicates that it could have further to go toward a more optimal point, and thus the step size α should be increased. However, if an update goes in the opposite direction as its previous update, it indicates that it could be oscillating, and thus the step size α should be decreased to lessen the magnitude of the oscillations. For example, Figure 5 shows the control decisions over time for a distributed controller. The left shows the power set point moving toward a more optimal point, where the right shows a controller that is oscillating.

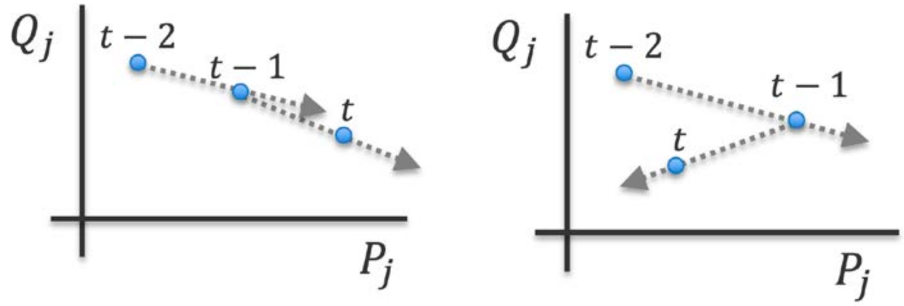

Figure 5. Distributed controller j: (left) indicates further to go, (right) indicates oscillations

We use the cosine similarity $s(\boldsymbol{x}_1, \boldsymbol{x}_2)$ as a metric that quantifies how much two $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ vectors go in the same or opposite directions:

$$s(\boldsymbol{x}_1, \boldsymbol{x}_2) \coloneqq \frac{\langle \boldsymbol{x}_1, \boldsymbol{x}_2 \rangle}{\|\boldsymbol{x}_1\| \|\boldsymbol{x}_2\|} \tag{6}$$

which gives a value between -1 and 1. Positive values indicate similar directions, while negative values indicate opposing directions. With the cosine similarity $s(\boldsymbol{x}_1, \boldsymbol{x}_2)$

However, we cannot simply use the cosine similarity $s(\boldsymbol{x}_1, \boldsymbol{x}_2)$ directly in RT-OPF DERMS if we want $\boldsymbol{x}_1$ to represent the update direction of the system at time t and $\boldsymbol{x}_2$ to represent the update direction of the system at time $t-1$. This is because the update direction of the system includes the updates of the dual and primal variables in Equations (1), (2), and (5), where the updates in Equations (1) and (2) are performed by the coordinator, and the updates in Equation (5) are performed by individual distributed controllers. Thus, we must either enhance the communication architecture of the RT-OPF DERMS to allow for all updates of the primal variables in Equation (5) to be sent back to the coordinator or allow the coordinator and distributed controllers to independently auto-tune an individualized local version of the step size α. We chose the latter option to preserve the communication architecture, and we further individualized the step size between voltage regulation and VPP because the conditions that cause variations in update directions can be quite different.

By individualizing the step sizes, the cosine similarity for the update directions at time t and $t-1$ can be defined as:

$$s_{volt}(t) \coloneqq s\big(\boldsymbol{d}_{volt}(t) - \boldsymbol{d}_{volt}(t-1), \boldsymbol{d}_{volt}(t-1) - \boldsymbol{d}_{volt}(t-2)\big) \tag{7a}$$

$$s_{VPP}(t) \coloneqq s\big(\boldsymbol{d}_{VPP}(t) - \boldsymbol{d}_{VPP}(t-1), \boldsymbol{d}_{VPP}(t-1) - \boldsymbol{d}_{VPP}(t-2)\big) \tag{7b}$$

$$s_{PQ,j}(t) := s\left(\begin{bmatrix} P_j(t) \\ Q_j(t) \end{bmatrix} - \begin{bmatrix} \hat{P}_j(t-1) \\ \hat{Q}_j(t-1) \end{bmatrix}, \begin{bmatrix} P_j(t-1) \\ Q_j(t-1) \end{bmatrix} - \begin{bmatrix} \hat{P}_j(t-2) \\ \hat{Q}_j(t-2) \end{bmatrix}\right) \tag{7c}$$

and implement the following threshold-based rule to increase, decrease, or keep constant each step size $\alpha_\theta(t) \colon \theta \in \{volt, VPP, (PQ, j)\}$:

$$\alpha_\theta(t) := \begin{cases} \overline{\gamma}\alpha_\theta(t-1), & if\ s_\theta(t) > \overline{s} \\ \alpha_\theta(t-1), & if\ \underline{s} \le s_\theta(t) \le \overline{s} \\ \underline{\gamma}_\theta \alpha_\theta(t-1), & if\ s_\theta(t) < \underline{s} \end{cases} \tag{8}$$

where $\underline{s}$ and $\overline{s}$ are lower and upper thresholds, respectively. The parameter $\overline{\gamma} > 1$ is the increase factor for all the step sizes, and $\underline{\gamma}_\theta$ is the individualized decrease factor. The reason to have the same increase factor $\overline{\gamma}$ used for all the DERMS components is to preserve the proportionality of the updates if the whole system accelerates by increasing the step sizes. On the other hand, having individualized decrease factors $\underline{\gamma}_\theta$ allows the engineer to set priorities between voltage regulation, VPP, and distributed controllers. Specifically, if oscillations occur in the system, the component with the smallest decrease factor $\underline{\gamma}_\theta$ will decrease its step size the fastest and thus have less dominance on the system to move toward its associated objective (i.e., voltage regulation, VPP, or local cost) (Comden, Wang, and Bernstein 2022).

Finally, to be explicit, we restate Equations (1), (2), and (5) with their individualized time-varying step sizes $\alpha_{volt}(t)$, $\alpha_{VPP}(t)$, and $\alpha_{PQ,j}(t) \colon \forall j$, respectively, as follows:

$$\underline{\boldsymbol{d}}_{volt}(t) := proj_{\mathbb{R}_+}\left\{\underline{\boldsymbol{d}}_{volt}(t-1) + \alpha_{volt}(t)\left(\underline{\boldsymbol{v}} - \hat{\boldsymbol{v}}(t-1) - \epsilon\underline{\boldsymbol{d}}_{volt}(t-1)\right)\right\} \tag{9a}$$

$$\overline{\boldsymbol{d}}_{volt}(t) := proj_{\mathbb{R}_+}\left\{\overline{\boldsymbol{d}}_{volt}(t-1) + \alpha_{volt}(t)\left(\hat{\boldsymbol{v}}(t-1) - \overline{\boldsymbol{v}} - \epsilon\overline{\boldsymbol{d}}_{volt}(t-1)\right)\right\} \tag{9b}$$

$$\underline{\boldsymbol{d}}_{VPP}(t) := proj_{\mathbb{R}_+}\left\{\underline{\boldsymbol{d}}_{VPP}(t-1) + \alpha_{VPP}(t)\left(\underline{\boldsymbol{P}}_0(t) - \hat{\boldsymbol{P}}_0(t-1) - \epsilon\underline{\boldsymbol{d}}_{VPP}(t-1)\right)\right\} \tag{10a}$$

$$\overline{\boldsymbol{d}}_{VPP}(t) := proj_{\mathbb{R}_+}\left\{\overline{\boldsymbol{d}}_{VPP}(t-1) + \alpha_{VPP}(t)\left(\hat{\boldsymbol{P}}_0(t-1) - \overline{\boldsymbol{P}}_0(t) - \epsilon\overline{\boldsymbol{d}}_{VPP}(t-1)\right)\right\} \tag{10b}$$

$$\begin{bmatrix} P_j(t) \\ Q_j(t) \end{bmatrix} := proj_{\mathcal{Y}_j(t-1)}\left\{\begin{bmatrix} \hat{P}_j(t-1) \\ \hat{Q}_j(t-1) \end{bmatrix} - \alpha_{PQ,j}(t)\left(\nabla_{P,Q}c_j\left(\hat{P}_j(t-1), \hat{Q}_j(t-1); t-1\right) + \begin{bmatrix} g_{P,j}(t) \\ g_{Q,j}(t) \end{bmatrix} + \nu\begin{bmatrix} \hat{P}_j(t-1) \\ \hat{Q}_j(t-1) \end{bmatrix}\right)\right\}. \tag{11}$$

2.1.3 List of Testing Scenarios of Automatic Tuning

In discussion with our industry partners, Utilidata and NVIDIA, we created four separate categories of scenarios to evaluate the designed automatic-tuning method under realistic scenarios:

1. Initial step sizes and convergence
2. Changes in VPP bounds
3. Solar power profiles
4. Changes in the tap ratio of a load tap changer (LTC).

The objective of the first category, Initial Step Sizes and Convergence, is to find a reasonable increase factor $\overline{\gamma}$ and decrease factors $\left(\underline{\gamma}_{volt}, \underline{\gamma}_{VPP}, \underline{\gamma}_{PQ}\right)$ that automatically tune the step sizes to reflect the priorities of a DERMS operator. The decrease factors put downward pressure on their respective step sizes; thus, the smaller the decrease factor, the smaller the priority is for that objective (e.g., voltage regulation, VPP, PV curtailment). This category is also to ensure that the automatic tuner will push the step sizes toward similar values, independent of their initial values. Finally, we have a testing scenario where all exogenous inputs are held constant to ensure that the control decisions converge toward a static optimal solution. The specific scenarios in this category are the following:

1. Initial step sizes and convergence:
    - 1.1. Sensitivity analysis of increase and decrease factors:
        - 1.1.1. Increase factor $\overline{\gamma}$
        - 1.1.2. Voltage regulation decrease factor $\underline{\gamma}_{\mathrm{volt}}$
        - 1.1.3. VPP decrease factor $\underline{\gamma}_{\mathrm{VPP}}$
        - 1.1.4. Distributed controllers decrease factor $\underline{\gamma}_{\mathrm{PQ}}$.
    - 1.2. Initial step sizes where only one starts far from converged values:
        - 1.2.1. Distributed controllers:
            - 1.2.1.1. Initial step sizes for all distributed controllers are low.
            - 1.2.1.2. Initial step sizes for all distributed controllers are high.
        - 1.2.2. Voltage regulation:
            - 1.2.2.1. Initial step size for voltage regulation is low.
            - 1.2.2.2. Initial step size for voltage regulation is high.
        - 1.2.3. VPP:
            - 1.2.3.1. Initial step size for VPP is low.
            - 1.2.3.2. Initial step size for VPP is high.
    - 1.3. All initial step sizes start far from converged values:
        - 1.3.1. All initial step sizes are low.
        - 1.3.2. All initial step sizes are high.
        - 1.3.3. Randomly generated initial step sizes.
    - 1.4. Static exogenous inputs to test for optimal convergence of control decisions.

The second category is focused on changes in VPP bounds. In all other categories, the VPP bounds are held constant. We tested how the automatically tuned step sizes performed under gradual linear changes, and under step changes in the VPP bounds. The primary performance characteristic is to see if the active power at the feeder head remains within the changing VPP bounds. The secondary performance characteristic is whether the voltage magnitudes stay within bounds, and the tertiary characteristic is whether the actions of the distributed controllers are smooth. Finally, we will be looking to how the step sizes change with the auto-tuner as the VPP bounds change. The specific scenarios are the following:

2. Changes in VPP bounds:
    - 2.1. Gradual linear changes
    - 2.2. Step changes.

The third category is focused on different solar power profiles. We tested scenarios that from ideal to extreme: with a large smooth solar profile, a smooth solar profile with large swings, and a profile with large sudden fluctuations. The priority of the performance characteristics is the same as in the second category, and the specific scenarios are labeled the following:

3. Solar power profiles:
    3.1. Smooth solar
    3.2. Large swings
    3.3. Large fluctuations.

The final category is focused on how the DERMS with auto-tuned step sizes reacts to changes in the system dynamics (e.g., the change of the tap ratio for the LTC at the substation). Usually, the LTC is controlled by ADMS. Because we did not have an ADMS in this validation platform, we simulated the change of LTC tap position as if it were changed by an ADMS. There are two main interests for this category. The first is to see how it would perform with an ADMS changing the tap ratio. The second comes from the internal information of the DERMS, specifically, the linearized power flow (see Equation (3)). We assume that the DERMS coordinator receives an admittance matrix of the feeder at the beginning and uses this information throughout its operation. When the tap ratio changes in the LTC, it changes the admittance matrix of the feeder. However, we do not assume that this information is relayed back to the DERMS coordinator, and so is operating based on the old admittance matrix, which constitutes larger model error than all other categories. Because the main objective of an ADMS is to keep the voltage magnitudes within bounds, that will be the primary performance characteristic we look at. The secondary characteristic will be the VPP. The specific labeled scenarios are:

4. Changes in the tap ratio of an LTC:
    4.1. Small changes in the tap position
    4.2. Large changes in the tap position.

2.1.4 Simulation Setup

HELICS is the tool that we used to evaluate the DERMS control performance in this project. We leveraged the platform from the ADMS test bed. The reason we chose HELICS is because it is a well-accepted tool for co-simulating grid models and various controls (e.g., DERMS, microgrid, DMS/ADMS) to evaluate the control performance of the tested system. For technical improvement, only a simulation is needed, and no hardware testing is involved. The HELICS framework developed for the technical improvement is shown in Figure 6. It has three agents, the OpenDSS agent, the DERMS coordinator agent, and the DERMS distributed local controllers' agent. The OpenDSS agent runs the a real-world distribution feeder with updated PV and load profiles, sends the measurements information to the DERMS coordinator; the DERMS coordinator receives the system-level measurements (e.g., feeder head power and selected nodes voltage), runs the algorithm with selected control objective(s) (voltage regulation and/VPP), sends the gradient values to the distributed local controllers; the DERMS distributed local controllers receive the gradient values and local measurements of the DER, and computes the PQ set points for each DER; and finally the PQ set points for DERs are applied in OpenDSS simulation.

Note that the distribution feeder has around 2,000 nodes, and 163 controllable nodes with PV and battery-connected. Each controllable node is a home. Therefore, we aim to control the behind-the-meter DERs to test the DERMS control performance in this technology commercialization fund project. Note that the PV generation and load profiles are collected from the field.

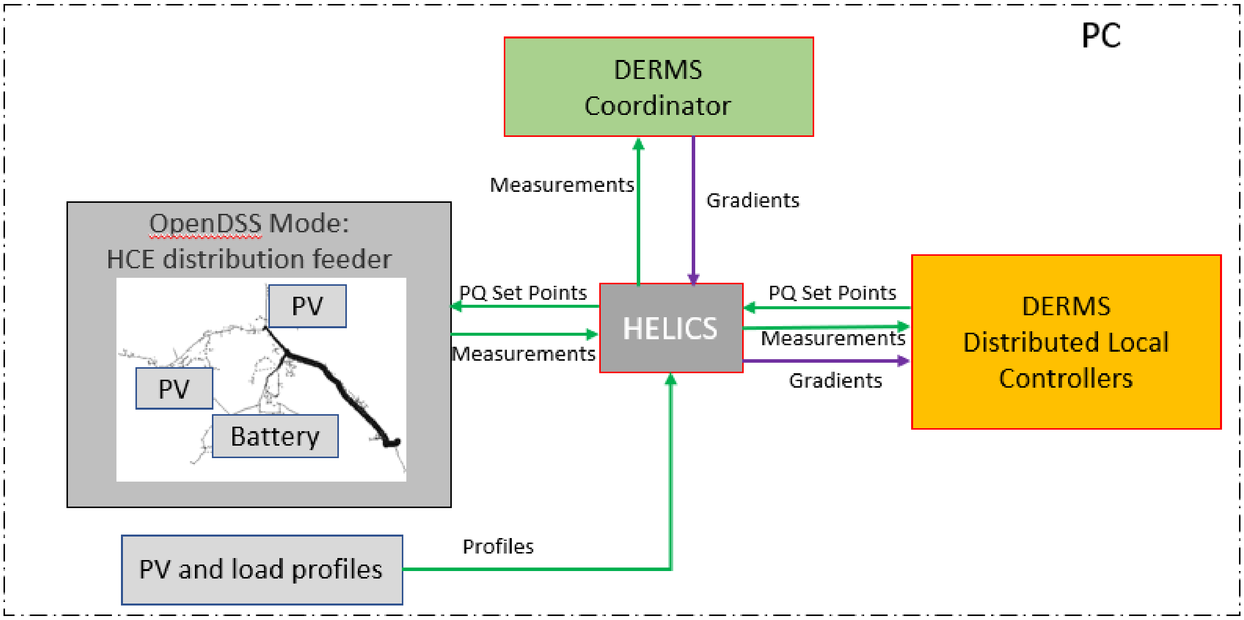

Figure 6. HELICS framework of the technology commercialization fund project

In these tests, the control objectives of the RT-OPF DERMS are voltage regulation and VPP. For the voltage regulation, the target voltage regulation limits are 1.03 for the upper limit and 0.95 for the lower limit. So, the voltages of all controllable nodes should be maintained within 0.95–1.03 p.u. For the VPP, the active power of each phase at the feeder head should be maintained within the band. The upper and lower limits are defined based on the power flow observed at the feeder head and the total available power for all the DERs (PV and battery). The simulations are for 2 to 6 hours, depending on the scenario being tested, and the scenarios will use a variety of trajectories of VPP bounds, solar profiles, and LTC tap positions. The load profiles update every 15 minutes, and the PV profiles update every 1 minute.

Although it may seem futile to use an auto-tuner that has its own parameters that need to be set, we show that it is less about tuning its own parameters and more about setting priorities for DERMS. We do this by using parameters for the auto-tuner that were set imprecisely based on assumed priorities.

The increase factor $\overline{\gamma}$ was set to 1.005, which allows a step size for any of the DERMS components to increase by roughly a magnitude every 15 minutes if the auto-tuner decides to continuously increase it. This is based on the frequency at which the DERMS components are making their updates, which is every 2 seconds. If the DERMS components were updating at different frequencies, then they could have different increase factors to preserve the proportionality in system action acceleration.

The decrease factors were set based on voltage regulation having higher priority than satisfying VPP bounds. In addition, the priority was limiting curtailment of PV generation via the distributed controllers between voltage regulation and VPP bounds. Thus, we set the descending order of the decrease factors based on those priorities to be $\underline{\gamma}_{\text{volt}} > \underline{\gamma}_{\text{PQ}} > \underline{\gamma}_{\text{VPP}}$. The decrease factor of the voltage regulation $\underline{\gamma}_{\text{volt}}$ as the highest priority was simply set to 0.995, which could decrease the step size by the same amount that could be increased by the auto-tuner. The decrease factor for the distributed controllers $\underline{\gamma}_{\text{PQ}}$ was set to 0.95, which has a 10-times proportional decrease in step size compared to voltage regulation. The decrease factor for VPP was set to 0.5, which has a 10-times proportional decrease in step size compared to the distributed controllers and 100-times proportional decrease compared to voltage regulation.

The lower threshold $\underline{\text{s}}$ and upper threshold $\overline{\text{s}}$ for the cosine similarity were set to 0 and 0.9 respectively, based on the suggested settings made by (Bernstein and Dall'Anese 2017). We will show in various scenarios that these rough parameter settings perform well in most realistic situations.

2.1.5 Simulation Results

2.1.5.1 Sensitivity Analysis of Increase and Decrease Factors (Scenario 1.1.X)

First, we explore the sensitivity that these increase factors and decrease factors have on the step sizes they are automatically tuning. We choose a day and time (April 3, 2019: 10 a.m.–12 p.m.) with a calm PV profile and with constant but active VPP bounds so that the step sizes will approximately converge. We analyzed the step sizes averaged over the last 15 minutes of the simulation as a representation of their converged values under perturbed increase and decrease factors. Under the base settings $\overline{\gamma} = 1.005$, $\underline{\gamma}_{volt} = 0.995$, $\underline{\gamma}_{VPP} = 0.5$, and $\underline{\gamma}_{PQ} = 0.95$, we have the following converged step sizes:

Table 1. Converged Step Size Values: Base Settings

α_{volt}	8.4×10^3
α_{VPP}	6.7
Average($\alpha_{PQ,j}$)	33

where the averaged step size among the 163 distributed controllers is reported. Next, we give the percentage changes in the converged step sizes when each of the factors is individually increased from the base settings.

Increasing the increase parameter $\overline{\gamma}$ from 1.005 to 1.01 gives:

Table 2. Change in Converged Step Size Values From Base Values: Scenario 1.1.1

Change in α_{volt}	+11%
Change in α_{VPP}	+22%
Average change in $\alpha_{PQ,j}$	+16%

Increasing the decrease factor $\underline{\gamma}_{volt}$ from 0.995 to 0.998 gives:

Table 3. Change in Converged Step Size Values From Base Values: Scenario 1.1.2

Change in α_{volt}	$+42\%$
Change in α_{VPP}	$+2.9\%$
Average change in $\alpha_{PQ,j}$	-13%

Increasing the decrease factor $\underline{\gamma}_{VPP}$ from 0.5 to 0.8 gives:

Table 4. Change in Converged Step Size Values From Base Values: Scenario 1.1.3

Change in α_{volt}	$+3.2\%$
Change in α_{VPP}	$+76\%$
Average change in $\alpha_{PQ,j}$	$+0.6\%$

Increasing the decrease factor $\underline{\gamma}_{PQ}$ from 0.95 to 0.98 gives:

Table 5. Change in Converged Step Size Values From Base Values: Scenario 1.1.4

Change in α_{volt}	-23%
Change in α_{VPP}	$+0.3\%$
Average change in $\alpha_{PQ,j}$	$+55\%$

The results of these perturbations show that the increase factor $\overline{\gamma}$ can be used to scale all the step sizes in a single direction, and changing the decrease factor $\underline{\gamma}_{VPP}$ can be used to set the priority of VPP with very little direct impact on the priority of the voltage regulation and the distributed controllers. However, changing either of the decrease factors $\underline{\gamma}_{volt}$ or $\underline{\gamma}_{PQ}$ to the priority of the associated objective has an inverse relationship on that of the other objective.

2.1.5.2 Initial Step Sizes Where Only One Starts Far From Converged Values (Scenario 1.2.X.Y)

All scenarios in the sensitivity analysis of the increase and decrease factors (Scenario 1.1.X) start with step sizes that are close to their converged values: $\alpha_{volt}(0) \coloneqq 5 \times 10^3$, $\alpha_{VPP}(0) \coloneqq 10$, $\alpha_{PQ,j}(0) \coloneqq 10$. In these test scenarios, we used the same base settings for the increase and decrease factors but shifted one initial step size for the objectives to be either extremely low (1/100$^{\text{th}}$) or extremely high (100 times) compared to the base initial values. We compared the converged values against the base scenario with initial step sizes close to their converged values.

Setting the initial step sizes for distributed controllers $\alpha_{PQ,j}(0)$ far from converged values gives:

Table 6. Change in Converged Step Size Values From Base Values: Scenarios 1.2.1.1 and 1.2.1.2

	$\boldsymbol{\alpha_{PQ,j}(0) \coloneqq 0.1}$	$\boldsymbol{\alpha_{PQ,j}(0) \coloneqq 1 \times 10^3}$
Change in α_{volt}	-3.6%	-30%
Change in α_{VPP}	-27%	$+15\%$
Average change in $\alpha_{PQ,j}$	$+1.3\%$	-1.3%

Setting the initial step sizes for voltage regulation $\alpha_{volt}(0)$ far from converged values gives:

Table 7. Change in Converged Step Size Values From Base Values: Scenarios 1.2.2.1 and 1.2.2.2

	$\boldsymbol{\alpha_{volt}(0) \coloneqq 50}$	$\boldsymbol{\alpha_{volt}(0) \coloneqq 5 \times 10^5}$
Change in α_{volt}	$+0.4\%$	-24%
Change in α_{VPP}	$+1.5\%$	-97%
Average change in $\alpha_{PQ,j}$	$+1.6\%$	-17%

Setting the initial step sizes for VPP $\alpha_{VPP}(0)$ far from converged values gives:

Table 8. Change in Converged Step Size Values From Base Values: Scenarios 1.2.3.1 and 1.2.3.2

	$\boldsymbol{\alpha_{VPP}(0) \coloneqq 0.1}$	$\boldsymbol{\alpha_{VPP}(0) \coloneqq 1 \times 10^3}$
Change in α_{volt}	$+1.0\%$	-23%
Change in α_{VPP}	$+19\%$	$+23\%$
Average change in $\alpha_{PQ,j}$	$+4.1\%$	$+6.8\%$

For most of the extreme initial step sizes starting two magnitudes low or high, the converged values are within 30% of those from the base scenario. One notable exception is the α_{VPP} ends 1/30th of that from the base scenario when $\alpha_{volt}(0)$ is set high. This is because with $\underline{\gamma}_{VPP} \coloneqq 0.5$, α_{VPP} is most sensitive to dropping fast when there were oscillations. This happened 30 minutes before the end of the 2-hour simulation, and the α_{VPP} was in the process of increasing before the simulation had ended (see Figure 7).

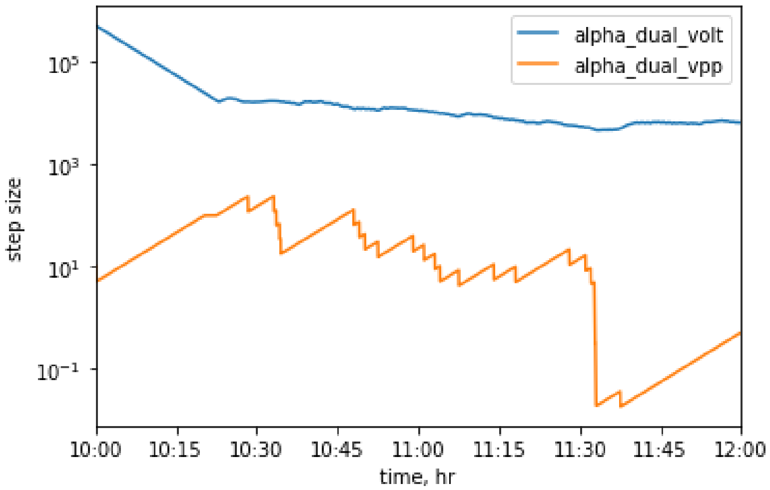

Figure 7. Step sizes for voltage regulation and VPP for Scenario 1.2.2.2

2.1.5.3 Initial Step Sizes With All Starting Far From Converged Values (Scenario 1.3.X)

In more realistic scenarios, the initial step sizes will all either be too low or too high, especially when first implementing RT-OPF DERMS. Also, during any random time in operation, the step sizes for the different DERMS components will likely be at different places. We tested these scenarios for the initial step sizes in the same way as Scenario 1.2.X.Y: all too low, all too high, or randomly distributed.

2.1.5.3.1 Scenario 1.3.1: Initial Step Sizes Set Far From Converged Values by Being Too Low

The initial step sizes were set to $\alpha_{PQ,j}(0) \coloneqq 0.1$, $\alpha_{volt}(0) \coloneqq 50$, and $\alpha_{VPP}(0) \coloneqq 0.1$, which gives the following converged step sizes compared to the base scenario:

Table 9. Change in Converged Step Size Values From Base Values: Scenario 1.3.1

Change in α_{volt}	-1.2%
Change in α_{VPP}	$+23\%$
Average change in $\alpha_{PQ,j}$	$+1.5\%$

These converged values, along with Figure 8 and Figure 9, show that even if the initial step sizes are all set very low, the auto-tuner quickly increases their sizes to a reasonable level and converges to them a value that is relatively close to the more optimistic scenario in which the initial values are close to the converged values.

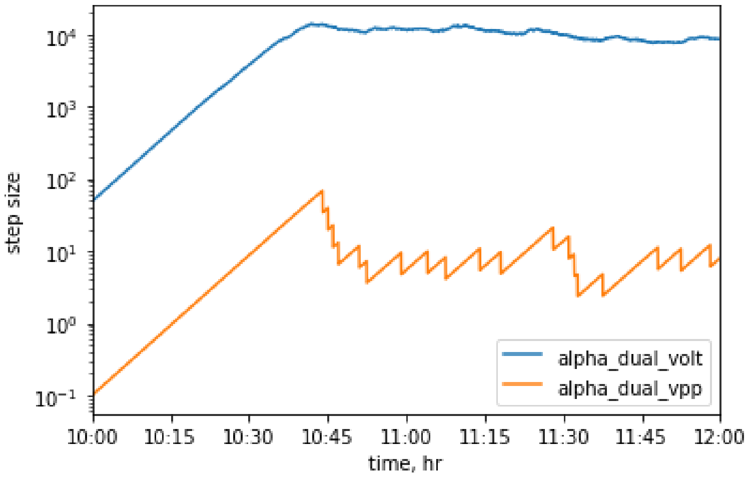

Figure 8. Step sizes for voltage regulation and VPP for Scenario 1.3.1

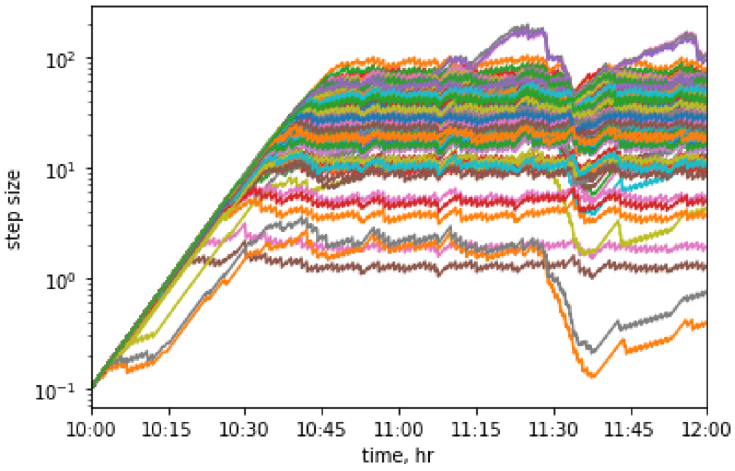

Figure 9. Step sizes for the distributed controllers for Scenario 1.3.1

2.1.5.3.2 Scenario 1.3.2: Initial Step Sizes Set Far From Converged Values by Being Too High

The initial step sizes were set to $\alpha_{volt}(0) \coloneqq 5 \times 10^5$, $\alpha_{VPP}(0) \coloneqq 1 \times 10^3$, and $\alpha_{PQ,j}(0) \coloneqq 1 \times 10^3$, which gives the following converged step sizes compared to the base scenario:

Table 10. Change in Converged Step Size Values From Base Values: Scenario 1.3.2

Change in α_{volt}	-31%
Change in α_{VPP}	$+5.8\%$
Average change in $\alpha_{PQ,j}$	-2.0%

Like Scenario 1.3.1, Figure 10 and Figure 11 show that if the initial step sizes are set too high, the auto-tuner can quickly bring them relatively close to the converged values of the base scenario. One notable difference comes from the fact that decrease factors are different so the initial drops in step sizes can be seen. And although the step sizes oscillate more than when they are set too low, they still level out in the end.

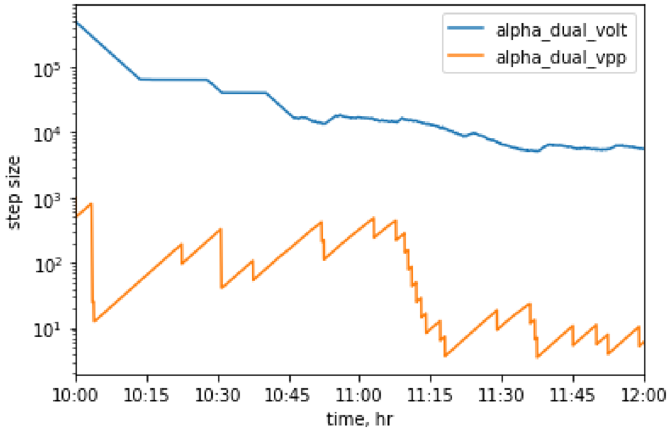

Figure 10. Step sizes for voltage regulation and VPP for Scenario 1.3.2

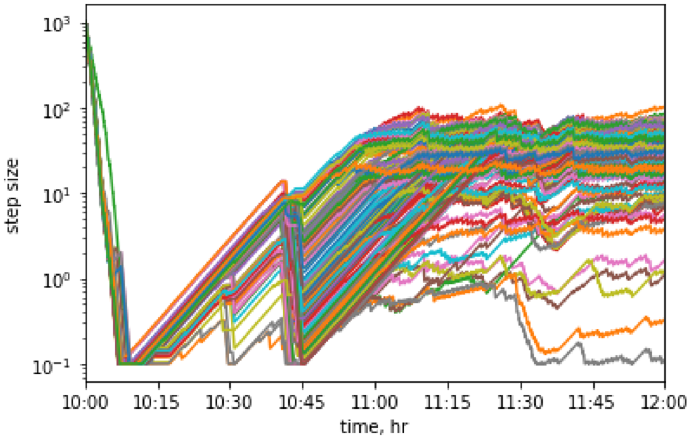

Figure 11. Step sizes for the distributed controllers for Scenario 1.3.2

2.1.5.3.3 Scenario 1.3.3: Initial Step Sizes at Random Values

The initial step sizes were randomly generated from a log-base-10 uniform distribution so that $\alpha_{volt}(0) \in [10^1, 10^4]$, $\alpha_{VPP}(0) \in [10^{-1}, 10^2]$, and $\alpha_{PQ,j}(0) \in [10^0, 10^3]$ where in this case $\alpha_{PQ,j}(0)$ was differentiated for each distributed controller. We ran three sample trials to show that although the initial step sizes were chosen within 3 magnitudes, their converged values were not far from the base scenario:

Table 11. Change in Converged Step Size Values From Base Values: Scenario 1.3.3.

	Trial #1	**Trial #2**	**Trial #3**
Change in α_{volt}	$+1.2\%$	-28%	-4.9%
Change in α_{VPP}	$+10\%$	$+1.6\%$	$+0.1\%$
Average change in $\alpha_{PQ,j}$	$+10\%$	$+6.8\%$	$+4.0\%$

2.1.5.4 Static Exogenous Inputs To Test for Optimal Convergence of Control Decisions (Scenario 1.4)

In this basic scenario, we set the PV, loads, and VPP bounds to be completely constant to verify that the control decisions from RT-OPF DERMS with step sizes that are being automatically tuned are continuously moving toward the minimum cost for a static system. This is important because a DERMS based on optimality should always be moving the system toward lower-cost or feasibility with respect to voltage magnitudes and VPP. This guarantee was already theoretically proved by (Bernstein and Dall'Anese 2017) for constant step sizes, and so we show it numerically for the automatically tuned DERMS by comparing it to the original (constant step size) DERMS. Figure 12 shows that the automatically tuned DERMS reaches the same total cost as DERMS with a constant step size set to 10. This gives us confidence that automatically tuning the step size does not break the theoretical guarantees given by the original RT-OPF DERMS.

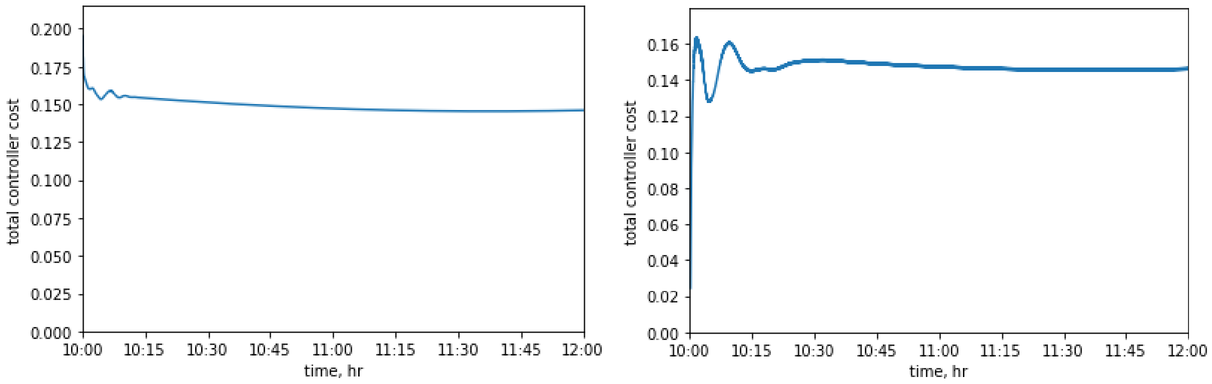

Figure 12. Total controller cost for (left) automatically tuned DERMS and (right) constant step size DERMS

2.1.5.5 Gradual Linear Changes in VPP Bounds (Scenario 2.1)

In this scenario, we took the base scenario and changed the VPP bounds to be time-varying with gradual linear changes to demonstrate that the auto-tuned DERMS can smoothly follow a gradually changing VPP set point for precise control. Figure 13 shows that the feeder head power can follow the gradually changing VPP bounds well. The small dips and spikes that occur every 15 minutes are from the load changing at the 15-minute granularity. Figure 14 shows that the voltage magnitudes also stayed within their prescribed bounds of 0.95 to 1.03 p.u. Figure 15 gives an example of the power injections from a distributed controller, which are smooth. The PV curtailment is 3.9% in this scenario.

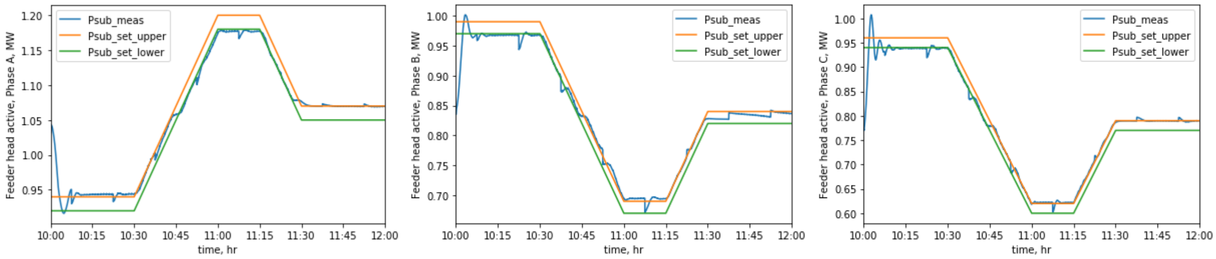

Figure 13. VPP bounds and feeder head powers of (left) Phase A, (middle) Phase B, and (right) Phase C for Scenario 2.1

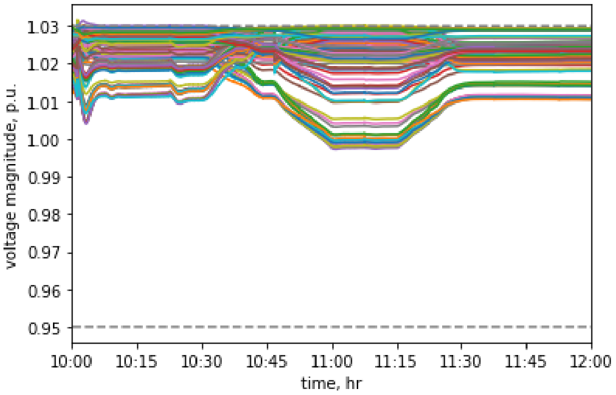

Figure 14. Voltage magnitudes for Scenario 2.1

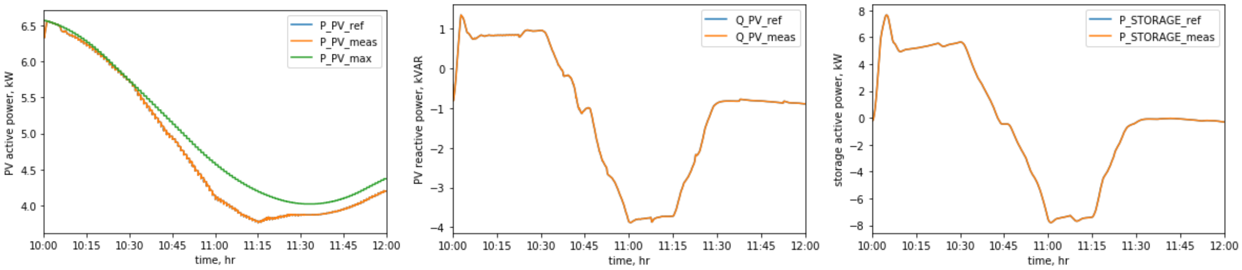

Figure 15. Distributed controller 53256.1 (left) PV active power, (middle) PV reactive power, and (right) storage active power injections for Scenario 2.1

2.1.5.6 Step Changes in VPP Bounds (Scenario 2.2)

This scenario was set up the same as Scenario 2.1, except that the VPP bounds were step changes to show how the DERMS can react to the sudden significant changes in the VPP set point. Figure 16 shows that the feeder head power can follow changes in VPP bounds very quickly and precisely, and Figure 17 shows that the voltage magnitudes also stayed within their prescribed bounds of 0.95 to 1.03 p.u. Figure 18 gives an example of the power injections from a distributed controller, which mirrors the step changes from the VPP bounds. The PV curtailment is 3.9% in this scenario.

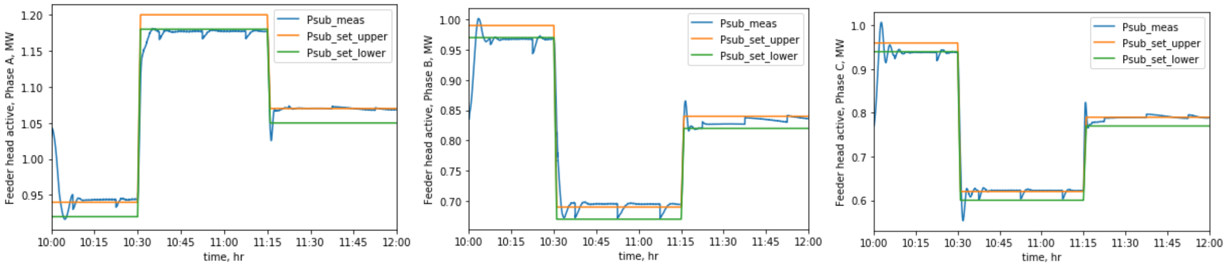

Figure 16. VPP bounds and feeder head powers of (left) Phase A, (middle) Phase B, and (right) Phase C for Scenario 2.2

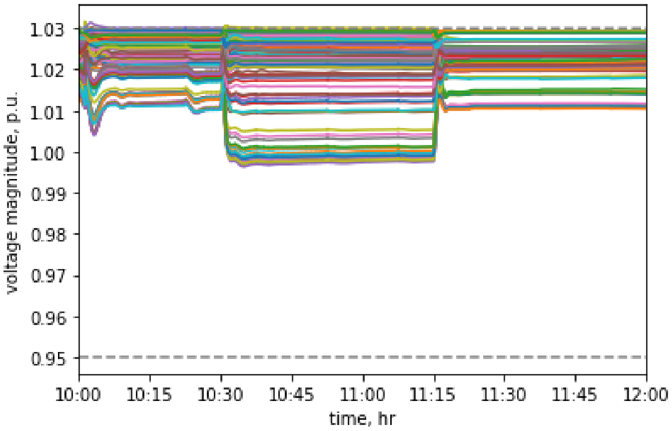

Figure 17. Voltage magnitudes for Scenario 2.2

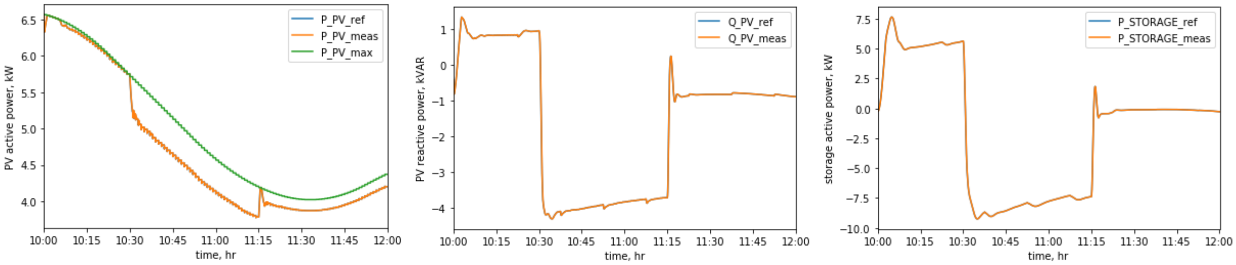

Figure 18. Distributed controller 53256.1 (left) PV active power, (middle) PV reactive power, and (right) storage active power injections for Scenario 2.2

2.1.5.7 Smooth Solar Power Profile (Scenario 3.1)

In this scenario, we held the VPP bounds constant but changed the day and time of the simulation so that the solar profile was at a high summer peak but smooth (July 4, 2019: 11 a.m.-–1 p.m.). Figure 19 gives the total solar profile used in the simulation. Figure 20 shows that the voltage magnitudes can be kept within bounds and are very smooth. Figure 21 gives the constant VPP bounds and feeder head powers for the three different phases. The VPP bounds were set to show different situations of the feeder head power almost completely unbounded (Phase A), completely bounded on one side (Phase B), and changing which VPP bound is active (Phase C). In all cases, the feeder head power is kept within bounds. The PV curtailment is 3.9% in this scenario.

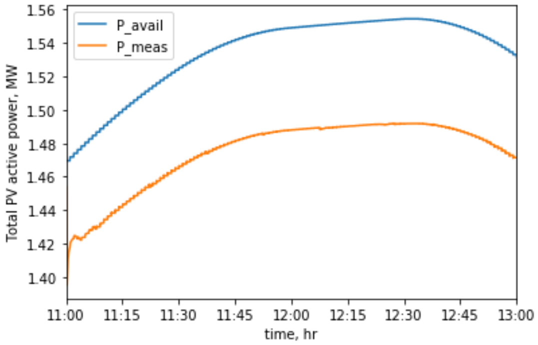

Figure 19. Total PV power available (blue) and used (orange) for Scenario 3.1

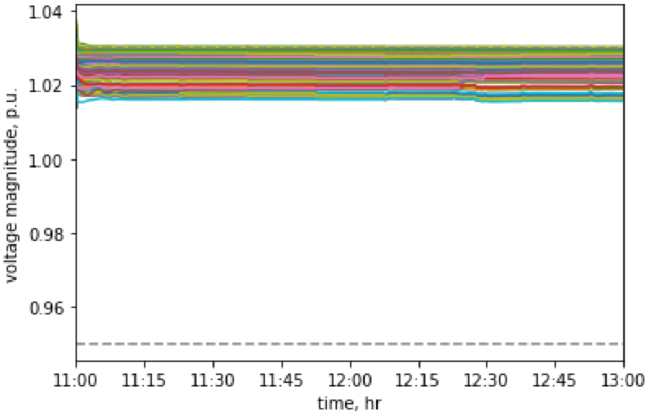

Figure 20. Voltage magnitudes for Scenario 3.1

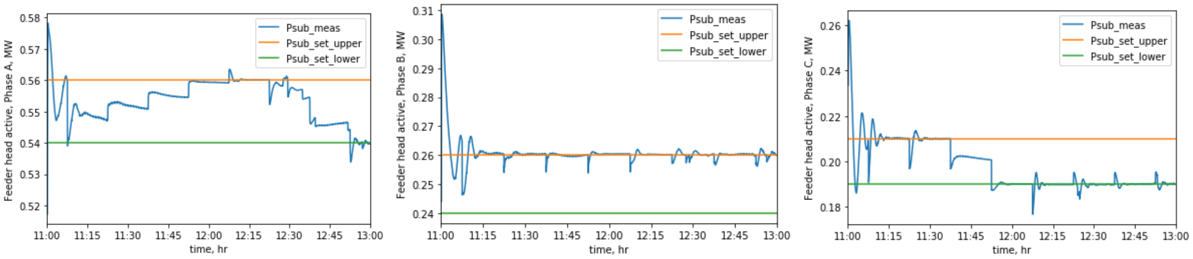

Figure 21. VPP bounds and feeder head powers of (left) Phase A, (middle) Phase B, and (right) Phase C for Scenario 3.1

2.1.5.8 Large Swings in the Solar Power Profile (Scenario 3.2)

Some days can see large dramatic swings in their solar profiles. In this scenario, we used a 6-hour time frame on June 4, 2019, from 9 a.m. to 3 p.m. with large swings that drop to as low as 2% of the peak (Figure 22). Figure 23 shows that the voltage magnitudes can be kept within bounds. Figure 24 shows the constant VPP bounds and feeder head powers. The swings in the PV power cause the feeder head power to be pushed outside of the VPP bounds, but in all the cases, the DERMS works to bring them close to the bounds. For example, in the first initial drop of 1.35 MW in PV power around 10 a.m., the feeder head powers only see a combined swing of 0.27 MW during that drop. The remaining drop in PV power gets absorbed by the energy storage. This can be seen in Figure 25, which gives an example of a distributed controller where the energy storage moves in the opposite direction of the PV to counteract the dramatic swings. We can also observe that the oscillations seen in the feeder head powers are not a result of oscillations in the decisions of the distributed controller, but it is instead of the fact that the energy storage is moving in opposite direction to the PV and is slightly delayed. The PV curtailment is 3.7% in this scenario.

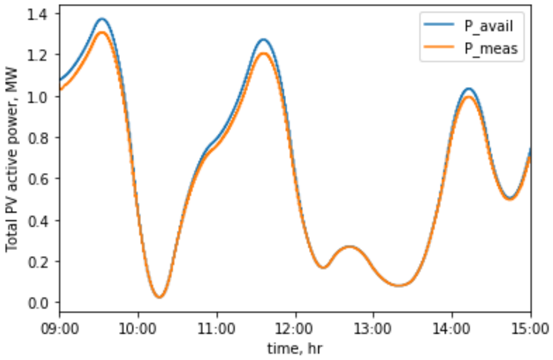

Figure 22. Total PV power available (blue) and used (orange) for Scenario 3.2

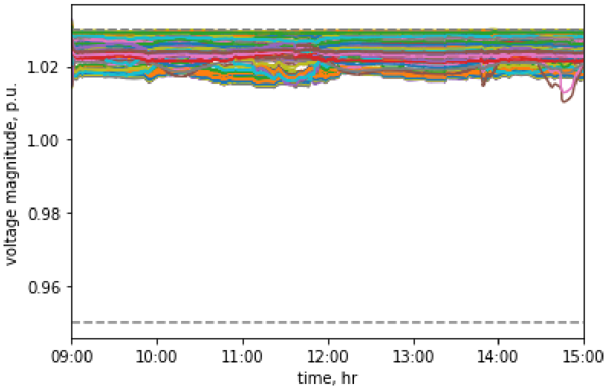

Figure 23. Voltage magnitudes for Scenario 3.2

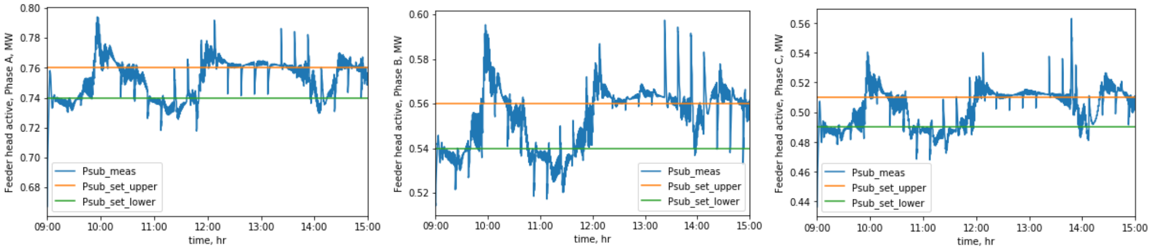

Figure 24. VPP bounds and feeder head powers of (left) Phase A, (middle) Phase B, and (right) Phase C for Scenario 3.2

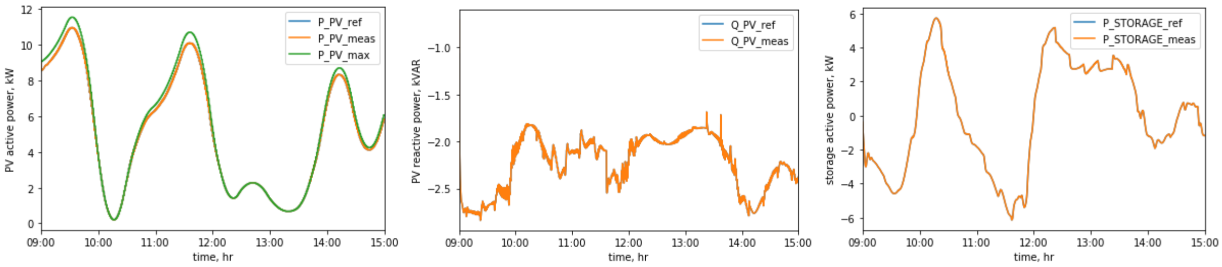

Figure 25. Distributed controller 53256.1 (left) PV active power, (middle) PV reactive power, and (right) storage active power injections for Scenario 3.2

2.1.5.9 Large Fluctuations in the Solar Power Profile (Scenario 3.3)

In the most extreme case, the solar power can have large fluctuations that are not smooth. We picked up a day with very fluctuated solar irradiance. We choose a 4-hour timeframe with fluctuations and combined with decreasing generation (Figure 26). Because of some of the difficulty for the VPP observed in Scenario 3.2, we increased the priority of the VPP by increasing $\underline{\gamma}_{VPP}$ from 0.5 to 0.995 and dropping $\underline{\gamma}_{PQ}$ from 0.95 to 0.8. It is reasonable to assume that a DERMS operator could receive a forecast that a partially sunny/cloudy day is coming and could make the adjustments in the auto-tuning parameters beforehand based on priorities. Figure 27 shows that the voltage magnitudes can be kept within bounds. Figure 28 shows the constant VPP bounds and feeder head powers. The large thin spikes in the feeder head powers coincide with the sudden large drops in PV power. This is because it takes one complete iteration of the DERMS to send an updated gradient signal to the distributed controller for the energy storage to absorb the immediately observed drop in PV power. Figure 29 shows the power injection for a distributed controller. Notice that the active power injection of the energy storage is almost a reflected image of the PV power, just slightly delayed. The PV curtailment is 3.7% in this scenario.

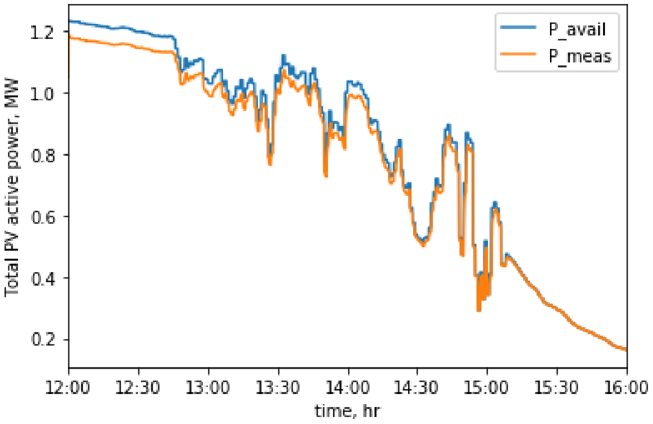

Figure 26. Total PV power available (blue) and used (orange) for Scenario 3.3

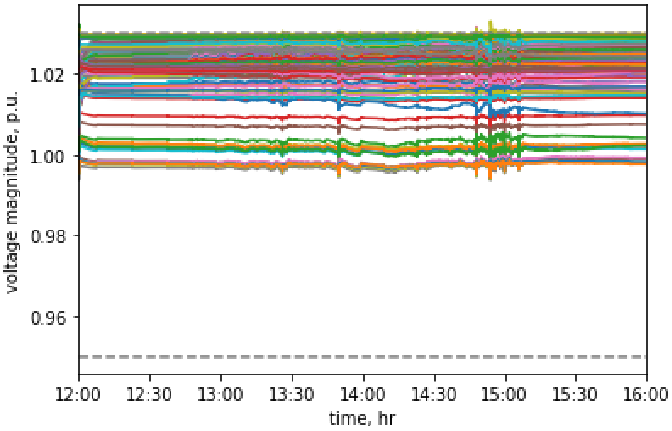

Figure 27. Voltage magnitudes for Scenario 3.3

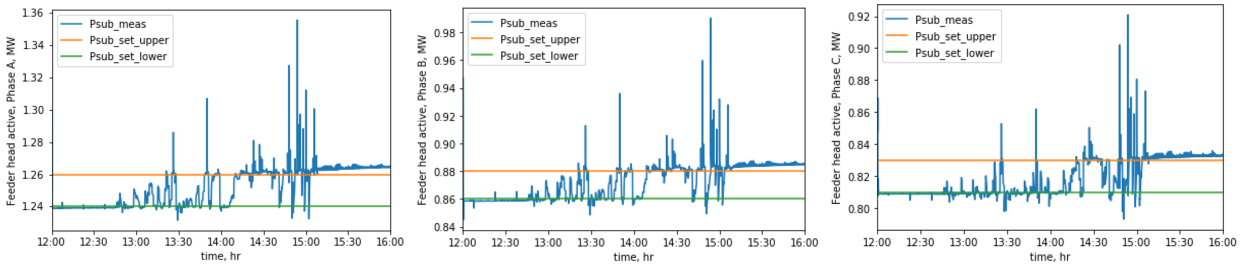

Figure 28. VPP bounds and feeder head powers of (left) Phase A, (middle) Phase B, and (right) Phase C for Scenario 3.3

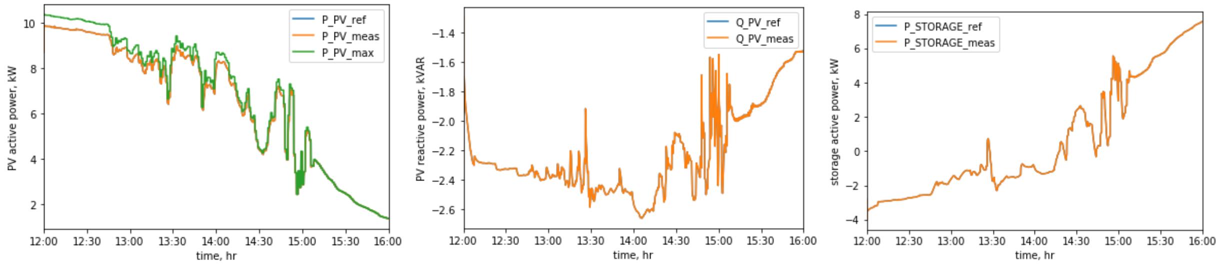

Figure 29. Distributed controller 53256.1 (left) PV active power, (middle) PV reactive power, and (right) storage active power injections for Scenario 3.3

2.1.5.10 Small Changes in the Tap Position (Scenario 4.1)

In this scenario, we tested the potential interaction between an ADMS, changing the tap position of an LTC located at the feeder head, and the RT-OPF DERMS. We used the base scenario with a smooth solar profile and constant VPP bounds but made small changes in the tap position, first down to reduce the voltage magnitude and then up to increase the voltage magnitude (Figure 30). Figure 31 shows the voltage magnitudes and the dual variables for the upper bound on the voltage magnitudes. Figure 32 shows the feeder head power and VPP bounds. The drop in the tap position from -1 to -3 at 10:30 a.m. causes the voltage magnitude to have a step size in the drop within bounds and is barely noticeable in the feeder head powers. However, at 11 a.m., the tap position is raised from -3 to +1; it immediately pushes many of the voltage magnitudes above the upper bound, causing oscillations in the voltage magnitudes and feeder head powers that are quickly dampened. This is because the sudden increase in the voltage magnitude makes the dual variables for the voltage magnitudes spike suddenly and send out gradients that signal the active powers to drop. Afterward, the feeder head powers being suddenly outside the bounds cause the gradients to be reversed and have them come back within bounds. This would be the case of an ADMS working against the DERMS (this is also known as the hunting effects). A way for a DERMS operator to lessen the impact of this kind of scenario would be to lower the priority of voltage regulation by lowering γ_{volt} since an ADMS would be sharing in that objective. The PV curtailment is 4.3% in this scenario.

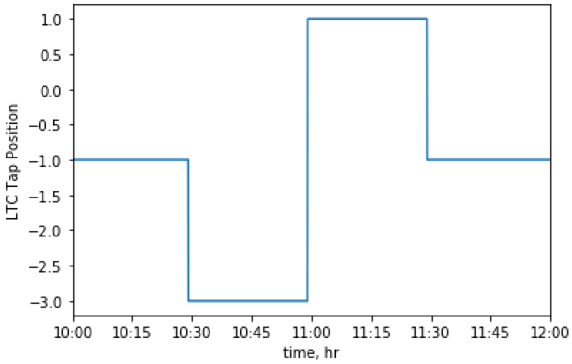

Figure 30. Tap position of the LTC for Scenario 4.1

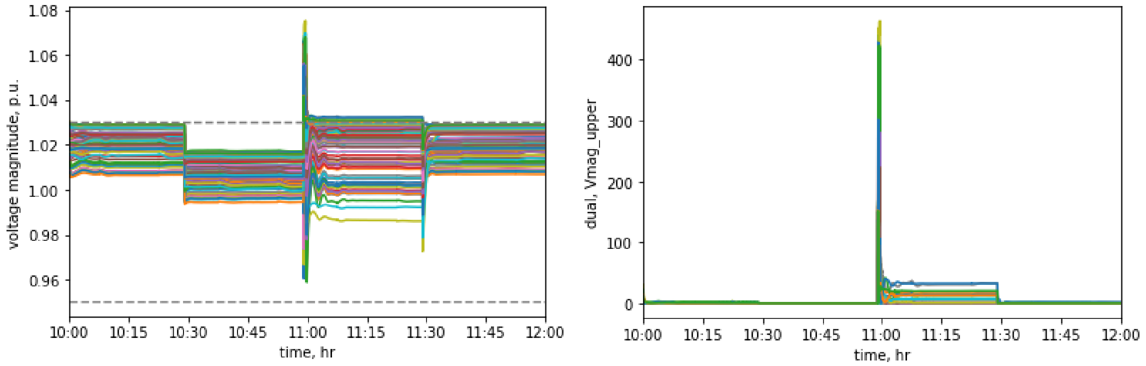

Figure 31. (Left) Voltage magnitudes for Scenario 4.1 and (right) dual variables for the upper bound on the voltage magnitudes

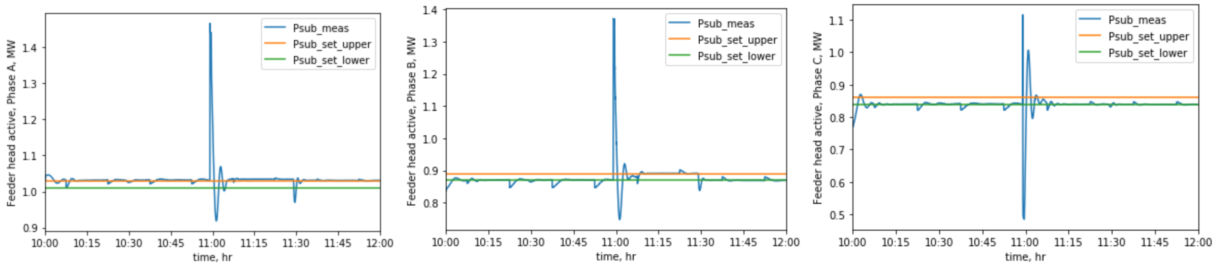

Figure 32. VPP bounds and feeder head powers of (left) Phase A, (middle) Phase B, and (right) Phase C for Scenario 4.1

2.1.5.11 Large Changes in the Tap Position (Scenario 4.2)

In this scenario, we repeated the same as Scenario 4.1, except the drop in the tap position goes much further down to -11 (Figure 33). Figure 34 shows the voltage magnitudes and the dual variables for the lower and upper and bounds on the voltage magnitudes. Figure 35 shows the feeder head power and VPP bounds. This scenario is very similar to Scenario 4.1, in that the observed dampened oscillations are not caused by the changes in the tap positions themselves, but how far this pushes the voltage magnitudes outside of the bounds. This can be seen for both the drop and rise in the tap position with the spikes in the dual variables for the voltage bounds. In fact, the initial spike in the feeder head powers at 11 a.m. are almost the same between the two scenarios (4.1 and 4.2), even though they are coming from two very different tap positions. Again, the solution for a DERMS operator in this case would be to lessen the priority of voltage regulation by DERMS if an ADMS is present. The PV curtailment is 4.0% in this scenario.

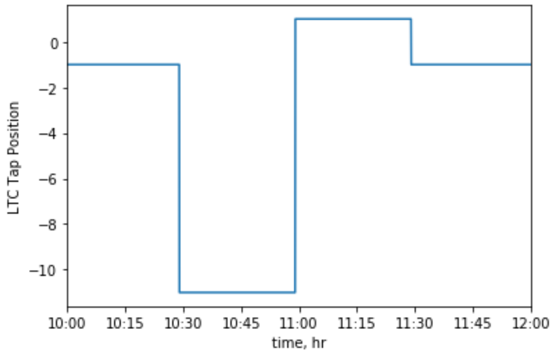

Figure 33. Tap position of the LTC for Scenario 4.2

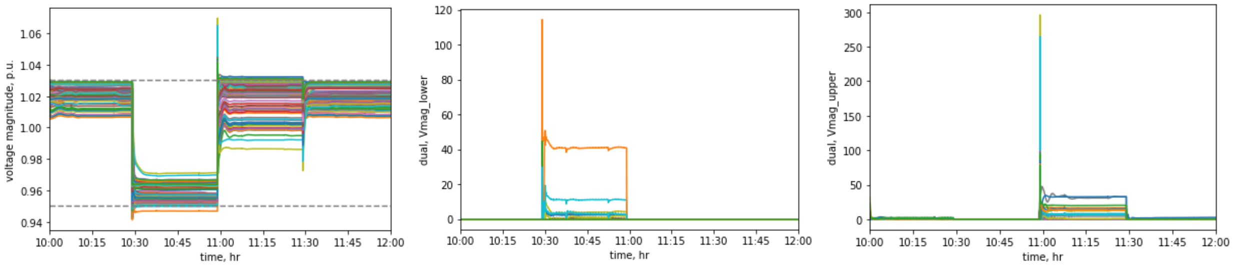

Figure 34. (Left) Voltage magnitudes for Scenario 4.1, (middle) dual variables for the lower bound and (right) dual variables for the upper bound on the voltage magnitudes

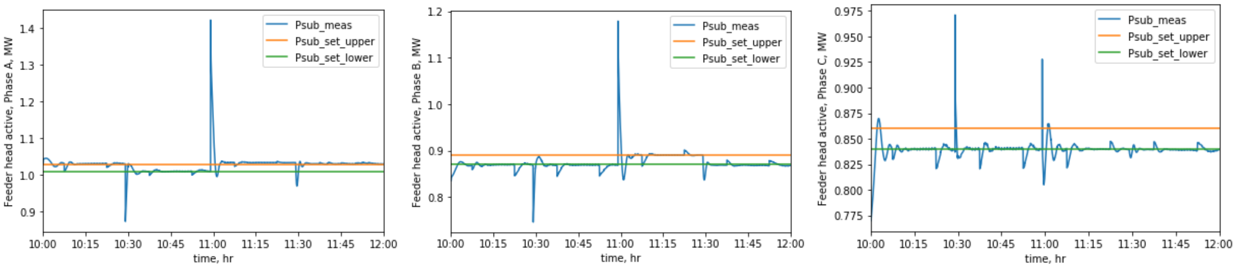

Figure 35. VPP bounds and feeder head powers of (left) Phase A, (middle) Phase B, and (right) Phase C for Scenario 4.2

2.1.6 Conclusion

We have shown in the first half of this project that the step sizes for an RT-OPF DERMS can be automatically tuned based on local variations in the gradients. This makes RT-OPF DERMS adaptable and can greatly lessen the time an engineer would need to get it up and running successfully on a system. The numerical simulations show that it is effective under many realistic scenarios that involve changes to the VPP set points, various solar profiles, and changes caused by an ADMS. Note that the simulation results of the last category suggest feasible coordinated control between ADMS and DERMS, which is an area that most utilities have no clue how to do and are still looking for solutions. With auto-tuning techniques, the NREL-developed RT-OPF DERMS is more user-friendly and efficient for deployment by industries, thus improving the maturity of the technology toward commercialization (Comden, Wang, and Bernstein 2023a).

2.2 Communication Issues Affecting the Performance of RT-OPF DERMS

The communication channel categories evaluated individually are:

1. Grid service measurements:
    a. Feeder head power measurements from the substation to the DERMS coordinator
    b. Voltage magnitude measurements from various nodes on the feeder to the DERMS coordinator.
2. Power injection direction signals from the DERMS coordinator to the DERMS local controllers
3. Power injection set points from DERMS local controllers to DERs, and power injection measurements from DERs to DERMS local controllers.

We display a diagram of the categories in Figure 36.

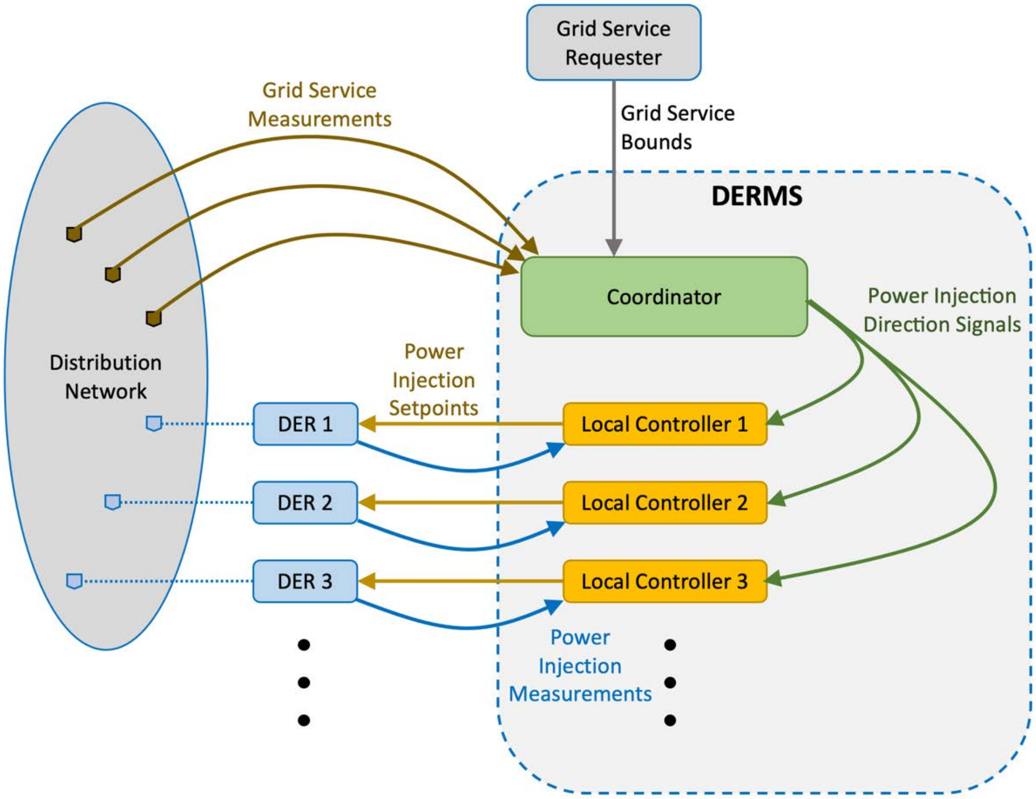

Figure 36. Information flow and communication channel categories of the DERMS

Based on the discussion with our industry partners (Utilidata and NVIDIA), we have the following communication aspects to evaluate.

Baseline: Fix the testing conditions, such as PV profile, load, VPP set points:

a. Execution time step: feeder agent 2 seconds, DERMS coordinator 5 seconds, local controllers 2 seconds
b. Smooth PV, VPP with one or two step change
c. Run simulation for 2 hours
d. With/without automatic tuning enabled.

S1: Evaluate the time step of coordinator and local controllers.

Requirement: Slow down the time step of testing agent from 5 seconds to minute(s) while fixing the time step of other agents and try to find the upper boundary of the testing agent. Note that the time step of local controllers is not studied because the time step of the coordinator dominates the performance of the RT-OPF DERMS, according to our previous study.

a. DERMS coordinator:
    - Gradually increase the time step
    - Lower bound of supervisory control and data acquisition 1–5 minutes.

S2: Evacuate the communication frequency.

Requirement: Slow down from 5 seconds to minute(s) and fix the time step of each agent. Try to find the upper boundary of communication frequency, and set coordinator time step of 1 minute for field deployment.

a. Feeder measurements to DERMS coordinator:
    - Gradually increase the time step
    - Lower bound of supervisory control and data acquisition 1 minute.
b. DERMS coordinator control signals to distributed local controllers
c. DER local controllers and DERs.

S3: Evaluate communication package loss.

Requirement: Try different probabilities 10% to 50% and aim to find the upper boundary of package loss of each communication path.

a. Feeder measurements to DERMS coordinator:
    - Two-dimension study: how many channels have package loss, what is the probability (how often) there is package loss.
b. DERMS coordinator control signals to distributed local controllers (two-dimension study)
c. DER local controllers and DERs (two-dimension study).

S4: Evaluate communication link failures.

Requirement: Choose random with probability of 0.1%–1% for several minutes at a time and occurrence chance is low.

a. Feeder measurements to DERMS coordinator
b. DERMS coordinator control signals to distributed local controllers
c. DER local controllers and DERs.

S5: Evaluate communication delays.

Requirement: Try imposing delay within 10–100 ms (may need to consider a few seconds delay), and test upper limit of the communication delays in each communication path.

a. Feeder measurements to DERMS coordinator (consider the actual delays in real-life implementation):
    - Voltage and power measurements (a few seconds delay).
b. DERMS coordinator control signals to distributed local controllers:
    - Randomly choose a few seconds delay up to minutes (need simulation to confirm what are delays making sense)
    - Two-dimension (how many of them have delays, how much are the delays).
c. DER local controllers and DERs:
    - Two-way communication and need to think of delays in both ways
    - Two-dimension (how many of them have delays, how much are the delays)
    - Delays are smaller than the one between coordinator and local controllers
    - Delays can be a few hundred microseconds to a few seconds.

2.2.1 Modified Setup for Evaluating Communication Issues

To evaluate how well RT-OPF DERMS can perform under different communication scenarios, we first had to modify our co-simulation setup in HELICS. The setup we had been using previously in the project (Figure 6): (1) simulated all the local controllers together in a single federate; and (2) used the publication/subscription method of sharing data between federates. These needed to be changed to allow for communication issues to be emulated.

The issue with (1) was that it did not allow for differentiation in communication delays and frequency among the local controllers. We handled (1) by splitting up the federate into separate federates for each local controller. The issue with (2) was that the publication/subscription method did not allow communication delays to be added into the simulation. We handled (2) by replacing the publication/subscription method with the endpoint method, which is a new feature within HELICS. The endpoint method allows for communication delays to data transfers, allows for data transfers to have origin and destination specifications, and mimics how data is sent over a communication network. Whereas the publication/subscription method treats data as a database that can be simply updated and queried at any time by any simulation federate without any delays. Figure 37 shows the current setup using the endpoint method.

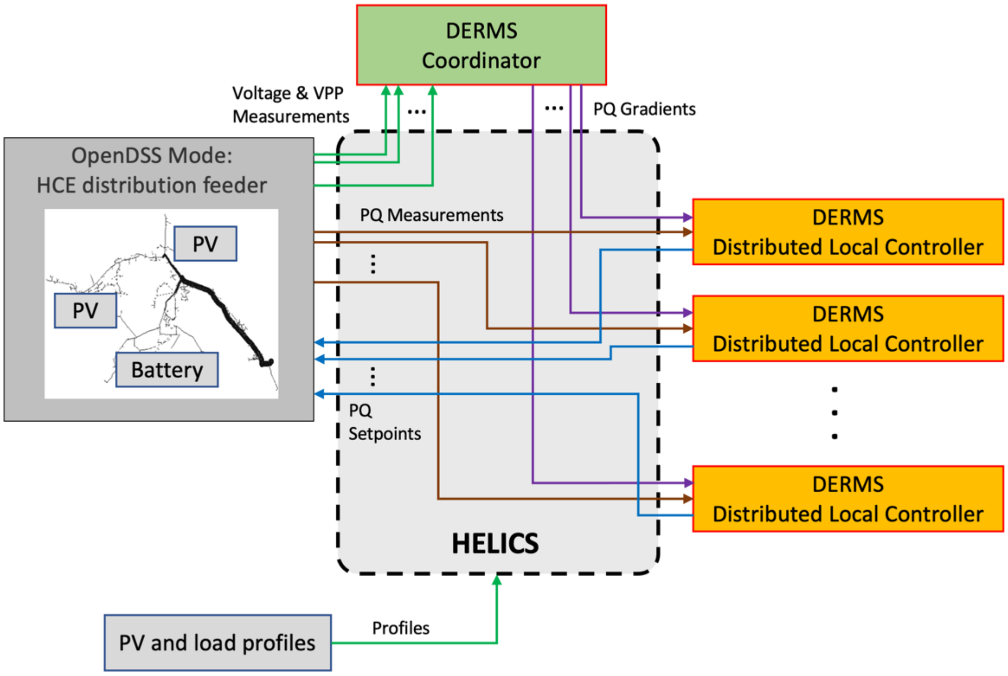

Figure 37. HELICS setup under the endpoint data-sharing method

(Comden, Wang, and Bernstein 2023d)

2.2.2 Communication Issue Evaluation

2.2.2.1 S1a: Upper Limit on the Update Frequency of the DERMS Coordinator

The frequency at which the data was communicated between the federates previously was every 2 seconds. However, in our discussions with Utilidata, a more realistic communication frequency between the feeder and the DERMS coordinator, and between the DERMS coordinator and the local controllers, is on the order of minutes. Thus, our first communication evaluation task was to find the upper limitation on the frequency of communication.

However, before that evaluation task could be done, the upper limit on the frequency of updating the DERMS coordinator needed to be evaluated, as that control element is central to collecting measurements from the distribution feeder and sending control signals to the distributed local controllers. We chose a 2-hour simulation scenario with a smooth PV profile and a large step change in the VPP set point so that we could see how quickly the DERMS responded to the changes while keeping the voltage magnitudes within their bounds. We tested this both with the manually tuned step size of 6, and automatic tuning with the increase factor set to 1.005 and all the decrease factors set to 0.9.

From our simulations under automatic step-size tuning, we found that the upper limit of the update frequency for the DERMS coordinator is 2 minutes. Figure 38 displays the VPP trajectories, and Figure 39 displays the voltage profiles. Under this scenario, it gradually brings the VPP trajectories to be within or very near their bounds and slowly brings the maximum voltage to within bounds. A slower update frequency (e.g., 3 minutes) makes the DERMS relatively ineffective—with respect to reaction time—at pushing the maximum voltage magnitude and the VPP trajectories of Phases A and B to be within bounds.

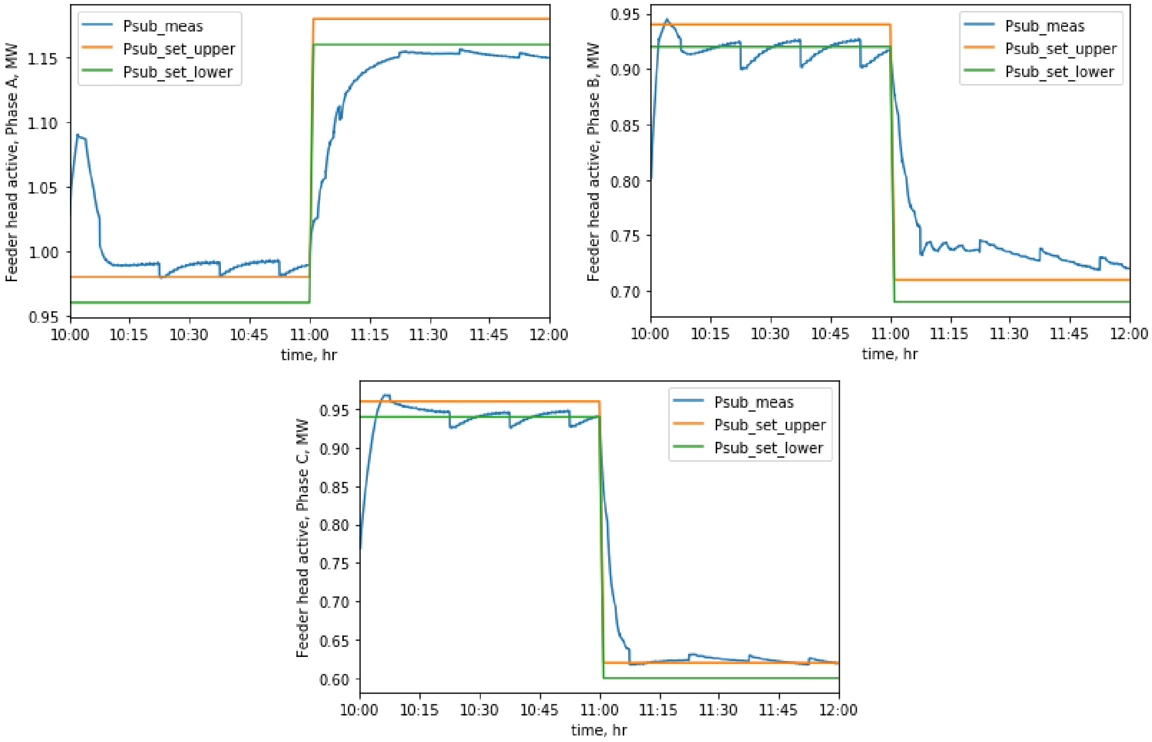

Figure 38. VPP trajectories for Phases A, B, and C, respectively, under automatic step-size tuning with the DERMS coordinator updating every 2 minutes

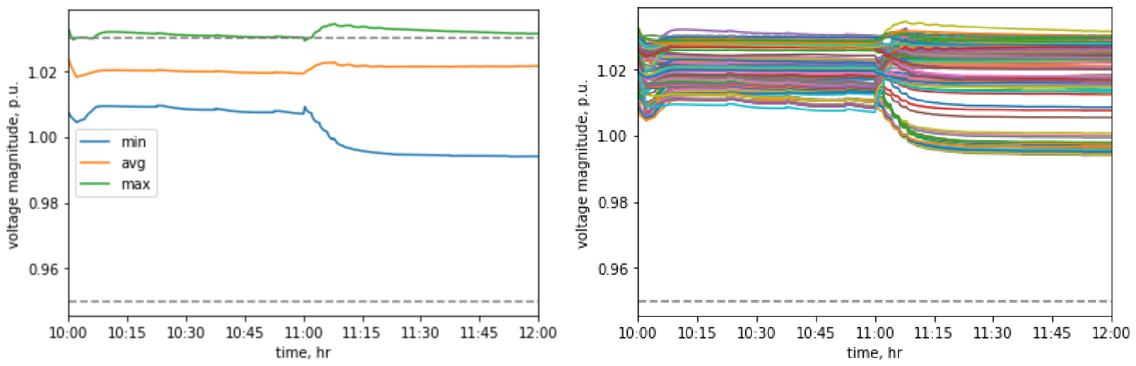

Figure 39. Voltage profiles under automatic step-size tuning with the DERMS coordinator updating every 2 minutes

Under manual step-size tuning, the upper limit of the update frequency for the DERMS coordinator was found to be 10 seconds, which is 1/12th that of the automatically tuned step sizes. This is because the manually tuned step size was set to 6, which made the DERMS react very slowly to voltage violations and the changing VPP bounds. The step size should be increased to make it react faster; however, increasing the step size above 6 makes at least one of the local controllers unstable and oscillate at a high frequency. The automatically tuned step sizes do not have this issue. Figure 40 displays the VPP trajectories, and Figure 41 displays the voltage profiles. Due to the upper limit of the update frequency of the manually tuned DERMS being on the order of seconds instead of minutes, the remaining evaluations will use the automatically tuned step sizes.

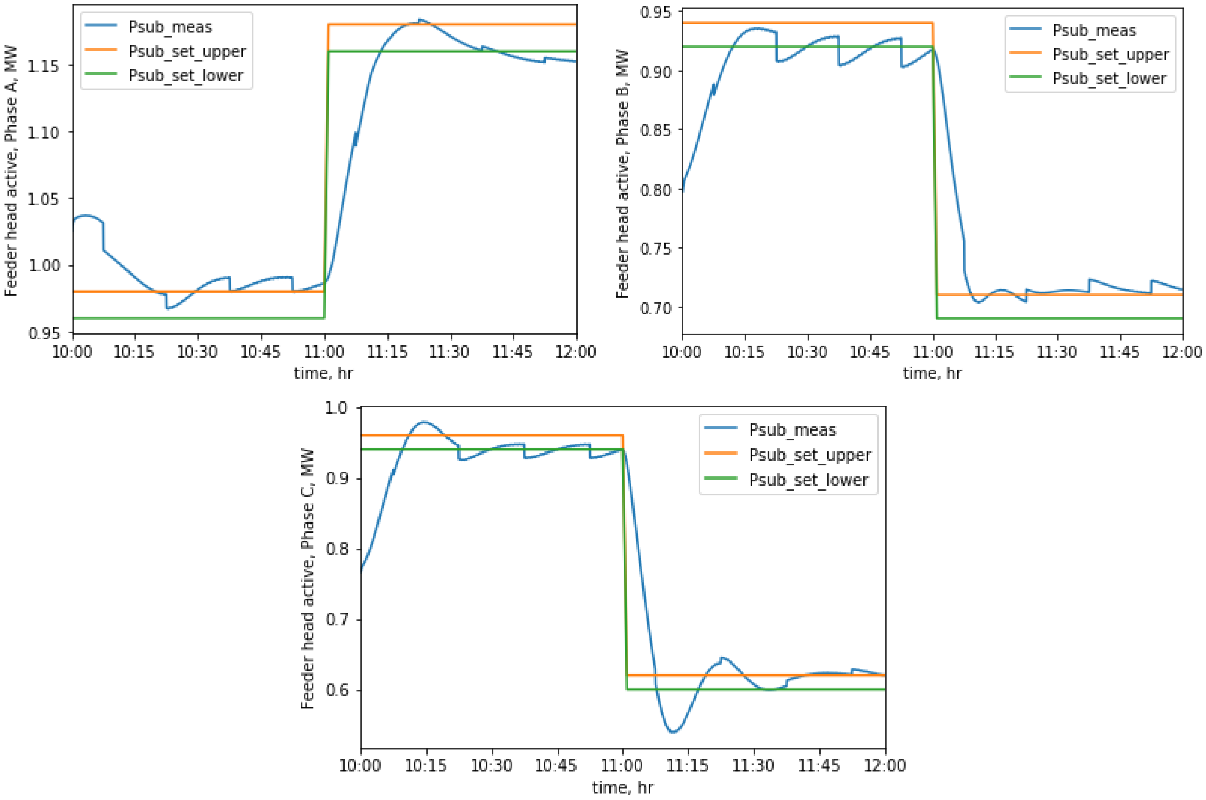

Figure 40. VPP trajectories for Phases A, B, and C, respectively, under manual step-size tuning with the DERMS coordinator updating every 10 seconds

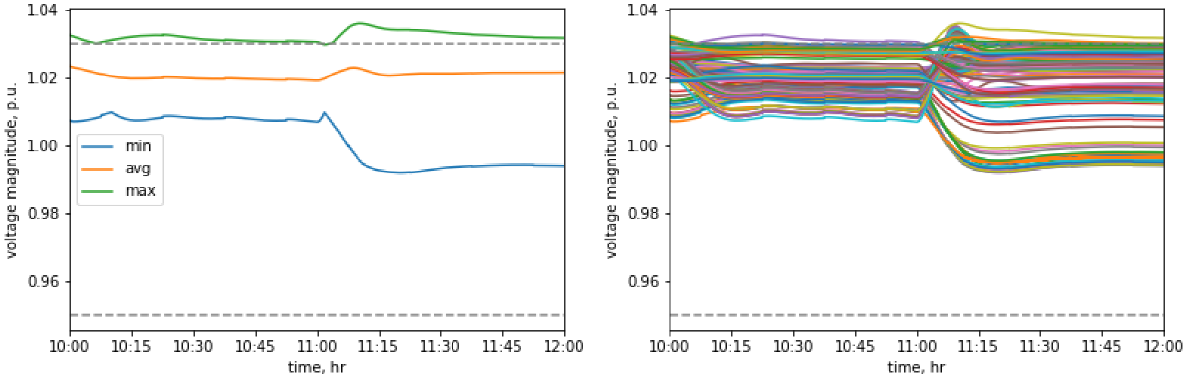

Figure 41. Voltage profiles under manual step-size tuning with the DERMS coordinator updating every 10 seconds

2.2.2.2 S2a: Upper Limit on the Voltage and VPP Measurement Frequency

The DERMS coordinator requires measurements of the voltage magnitudes from various locations on the feeder and powers at the feeder head (see green arrows in Figure 37). Our simulations were originally set to have the coordinator receive these measurements every 2 seconds. However, real-time information from supervisory control and data acquisition is typically given in the order of minutes and not seconds. In this set of simulations, we increased the voltage and VPP measurement frequency until the DERMS was no longer functional. The update frequency of the coordinator was set to 1 minute.

The upper limit on the voltage and VPP measurement frequency was found to be 2 minutes. Figure 42 displays the VPP trajectories, and Figure 43 displays the voltage profiles. The results are very similar to the upper limit of the update frequency for the DERMS coordinator (see Figure 38 and Figure 39).

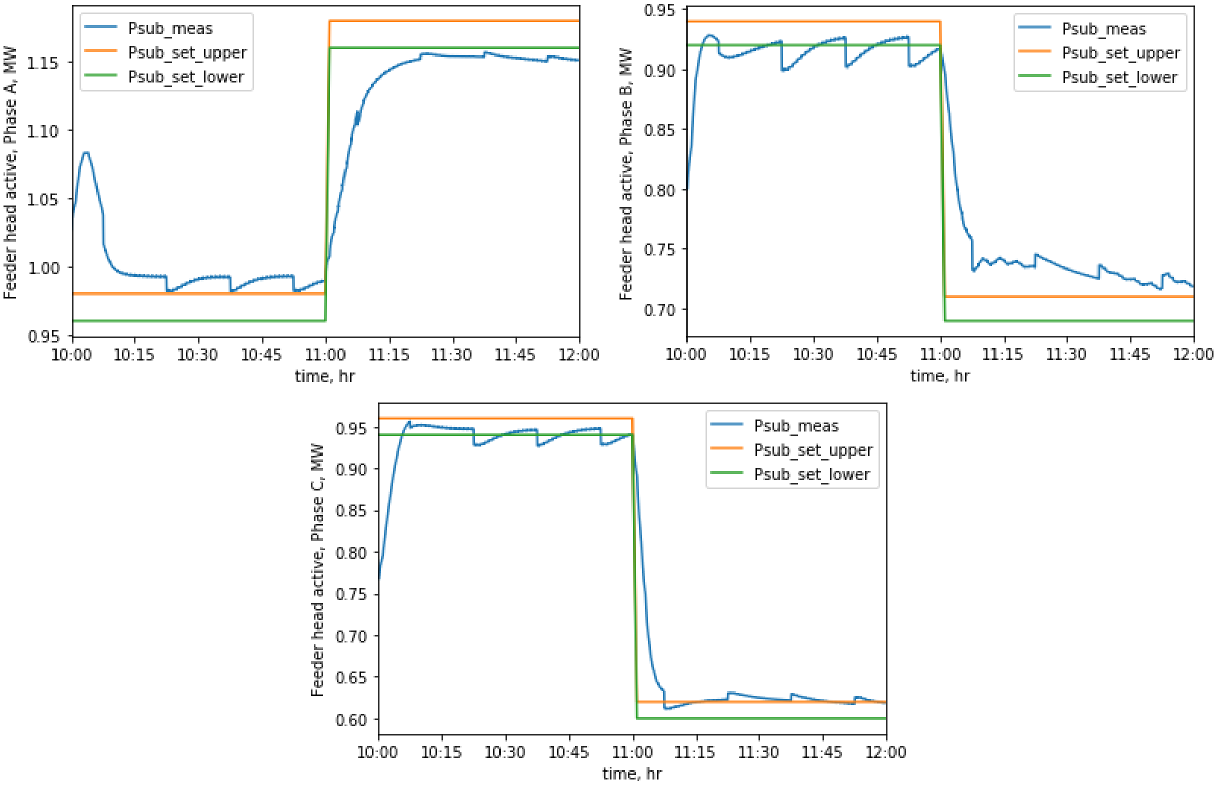

Figure 42. VPP trajectories for Phases A, B, and C, respectively, under automatic step-size tuning with the voltage magnitudes and feeder head power measured every 2 minutes

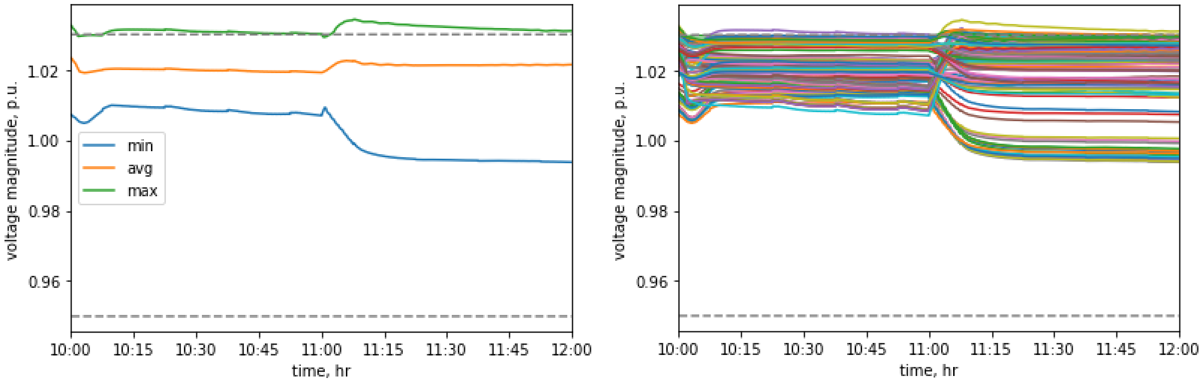

Figure 43. Voltage profiles under automatic step-size tuning with the voltage magnitudes and feeder head power measured every 2 minutes

2.2.2.3 S2b: Upper Limit on the Frequency That the DERMS Coordinator Sends Gradients to the Local Controllers

Additionally, we tested the upper limit on the frequency that the DERMS coordinator sends the PQ gradients to the local controllers (see the purple arrows in Figure 37) because they are sent over a wide area communication network. Our simulations were originally set to have them sent every 2 seconds. The update frequency of the coordinator is set to 1 minute.

The upper limit on the frequency that the PQ gradients are sent to the local controllers was found to be 2 minutes. Figure 44 displays the VPP trajectories and Figure 45 displays the voltage profiles. The results are very similar to the upper limit of the update frequency for the DERMS coordinator (Figure 38 and Figure 39).

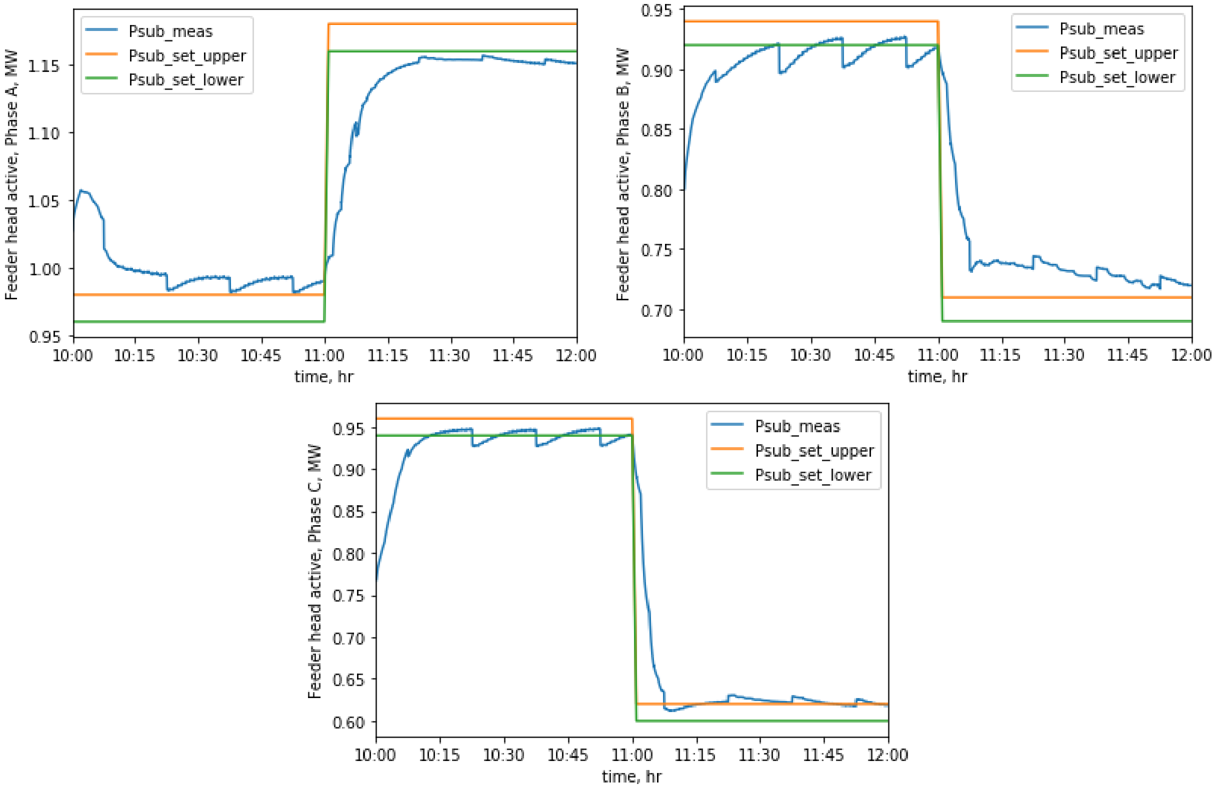

Figure 44. VPP trajectories for Phases A, B, and C, respectively, under automatic step-size tuning with the PQ gradients sent to the local controllers from the coordinator every 2 minutes

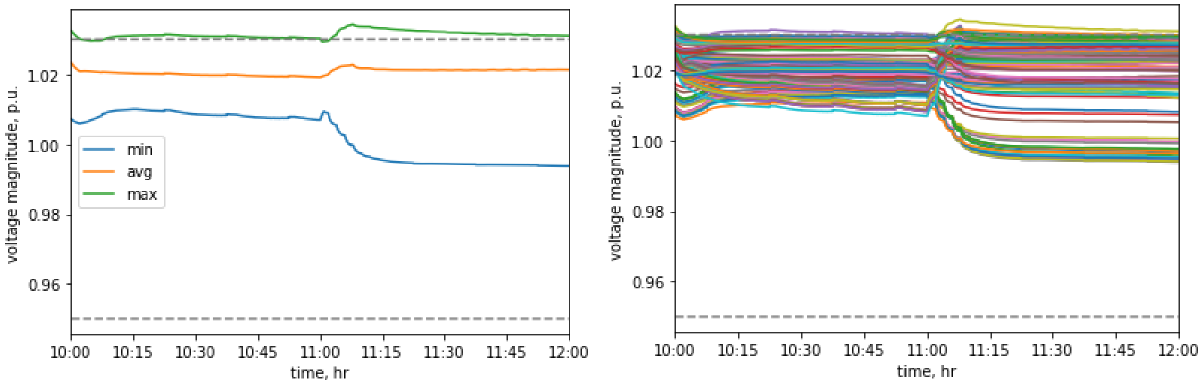

Figure 45. Voltage profiles under automatic step-size tuning with the PQ gradients sent to the local controllers from the coordinator every 2 minutes

2.2.2.4 S2c: Upper Limit on the Frequency That a Distributed Local Controller Communicates With its Associated DERs

The last communication frequency loop that needed to be evaluated was the bidirectional communication between the local controller and its associated DER(s) that it controls (see the brown and blue arrows in Figure 37). The default settings were originally set so that each local controller and DER communicated every 2 seconds. This is reasonable if the local controller and the DER are located close to each other; however, we wanted to test the upper limit on the frequency of communication. The update frequency of the coordinator was set to 1 minute.

Initially, we changed the frequency of bidirectional communication to 3 seconds while keeping the frequency at which a local controller updates its PQ set points to 2 seconds. However, this made the DERMS fail in bringing the voltages or the VPP to be within their bounds. This is because the update step in the local controller uses the measured PQ value as a reference point from which to adjust (see Equation (11)). If the PQ measurement does not arrive in time for the next update step, it could use the previously received value; however, then this means that the next PQ set point may not be moving in the correct direction in relation to the currently implemented set point. This mismatch breaks the logic of how the local controller should be directing the DER.

Fortunately, there was a simple solution to this issue—matching the frequency of the local controller update with the frequency that it communicates with the DER(s) it controls. By doing this, we successfully pushed the frequency above 2 seconds, to an upper limit of 1 minute. Moving forward, this indicates that the local controller should update the set point it sends to the DER only after it receives an updated measured PQ value from its DER(s). Figure 46 displays the VPP trajectories, and Figure 47 displays the voltage profiles. Notice that they look very different than in our previous simulations in two major ways. The first is that the VPP has larger oscillations. And the second is that the degradation of the VPP grid service is what limits the communication and update frequency from being pushed up higher, whereas in the previous simulations, it was the voltage regulation grid service that degraded first.

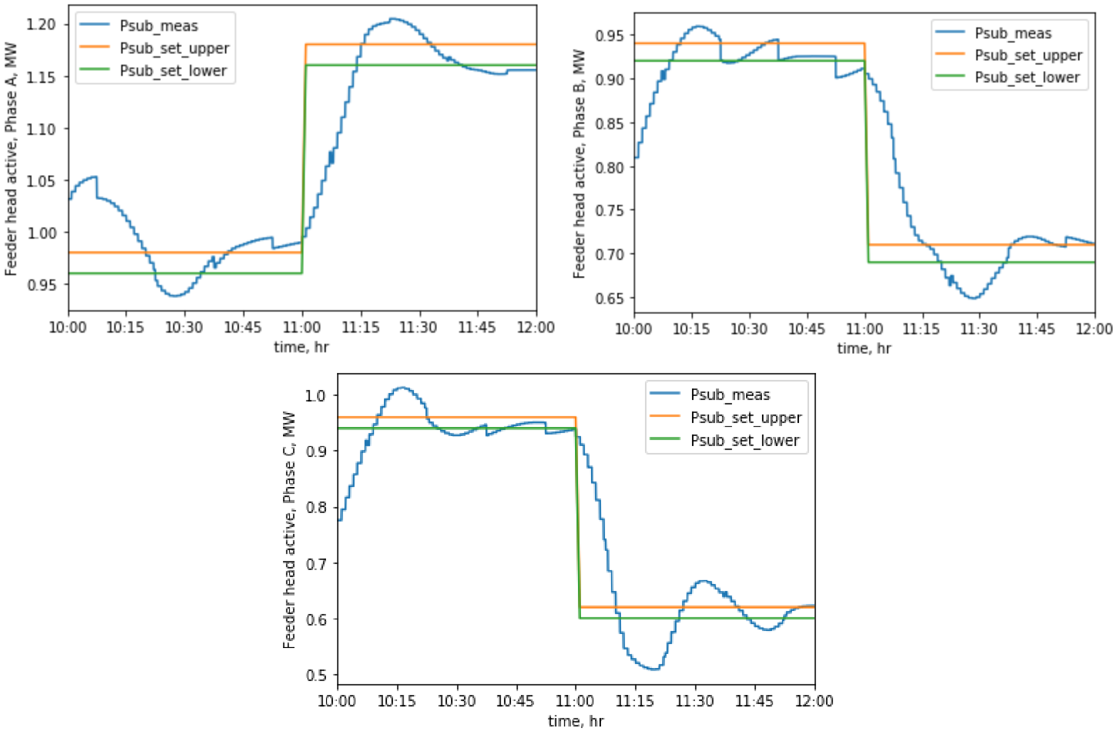

Figure 46. VPP trajectories for Phases A, B, and C, respectively, under automatic step-size tuning with the local controllers communicating with their associated DER(s) and updating their set points every 1 minute

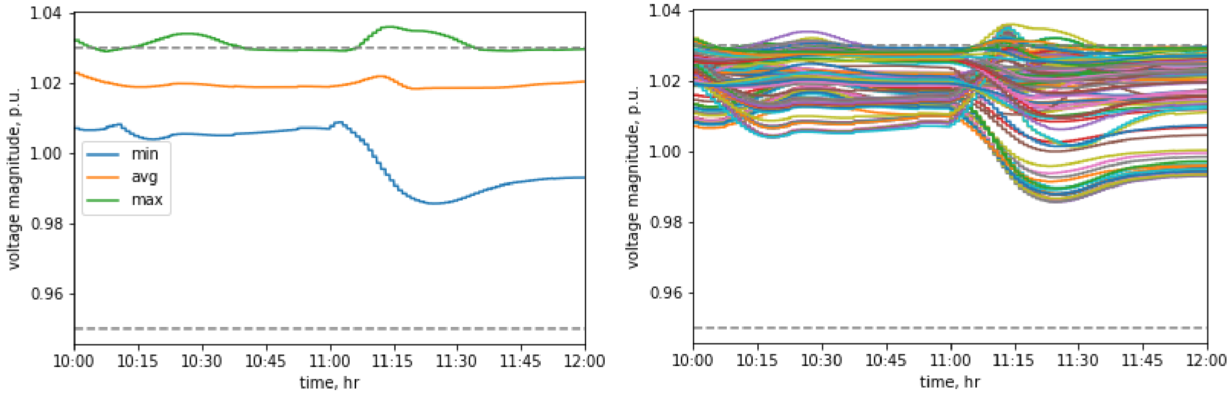

Figure 47. Voltage profiles under automatic step-size tuning with the local controllers communicating with their associated DER(s) and updating their set points every 1 minute

2.2.2.5 S3: Random Packet Drop

When utilizing communication networks to pass messages between different DERMS components, messages can occasionally drop and not make it to their destination. These drops usually happen without warning and are one-off events. Thus, we modeled this phenomenon as a Bernoulli process in which every time a message is sent, there is a specific probability that the message will not be delivered to its destination. We evaluated each category of communication channel separately. For each category, we evaluated the upper bound on the probability of packet loss that will keep the DERMS functional. We varied the number of channels in the category set to have random pack loss to find the upper bound on the probability as a function of the number of channels. All other channels are assumed to have no random packet loss.

Within HELICS, the random packet drop for the grid service measurements (feeder head power measurements and voltage magnitudes) is simulated by adding a native filter to the sending endpoint of the communication channels from the feeder to the DERMS coordinator (the origin of the green arrows in Figure 37). The random packet drops for the power injection direction signals, sending PQ set points, and receiving PQ measurements were simulated within local controller agents. The time step granularity of the grid service measurement communication, DERMS coordinator update, and power injection direction signal communication was set to 1 minute to match the supervisory control and data acquisition limitations as determined in the work for communication issues S1 and S2. The time step granularity of the communication between the local controllers and the DERs was set to 10 seconds.

First, we evaluated the upper bound on the probability of packet drop versus the number of feeder head power measurement communication channels, separated by phases, with a random packet drop (see Figure 48). Because there are only three channels, we could explore all phase permutations. When all feeder head power measurement channels have at most 0.4 probability of packet drop, then the DERMS is still functional. Phase C defines the lower envelope, which indicates that it has the most effect on the loss of functionality.

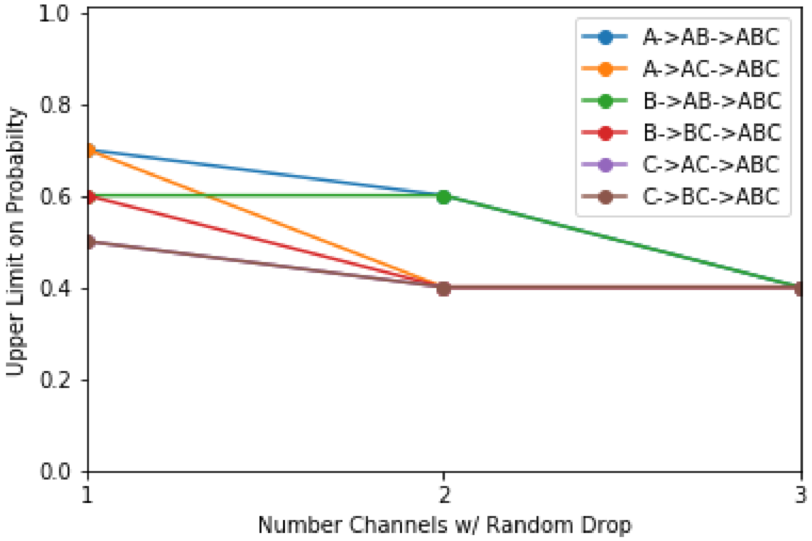

Figure 48. Feeder head power measurements: upper bound on the probability of packet drop versus number of channels/phases

The legend shows which permutations of phases have the possibility of packet drop.

Second, we evaluated the upper bound on the probability of packet drop versus the number of voltage magnitude measurement communication channels from measurement nodes with a random packet drop (Figure 49). When all voltage magnitude measurement channels have at most 0.5 probability of packet drop, then the DERMS is still functional. There is a threshold at 143 (88%) channels where the upper bound on the probability is 0.9. After this threshold, the upper bound on the probability drops significantly. This shows that the DERMS can be functional if at least 12% of the channels were on a reliable communication network with regards to the packet drop, and the rest of them could be on an unreliable network.

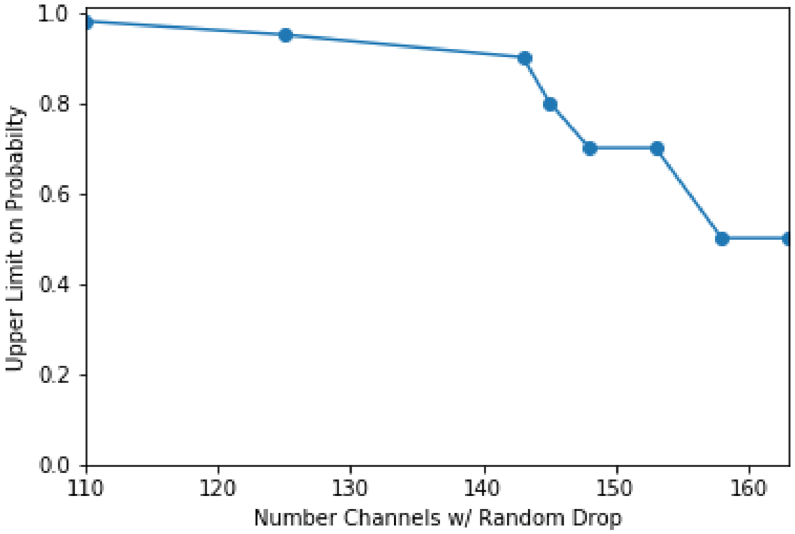

Figure 49. Voltage magnitude measurements: upper bound on the probability of packet drop versus number of channels/nodes

Third, we evaluated the upper bound on the probability of a packet drop versus the number of power injection direction signal communication channels with random packet drops (Figure 50). These are the internal communication channels of the DERMS from the coordinator to the distributed local controllers. When all power injection direction signal channels have at most 0.28 probability of packet drop, then the DERMS is still functional. There is a threshold at 117 (72%) channels where the upper bound on the probability is 0.83. After this threshold, the upper bound on the probability drops significantly. This shows that the DERMS can be functional if at least 28% of the channels are on a reliable communication network with regards to packet drop, and the rest of the channels could be on an unreliable network.

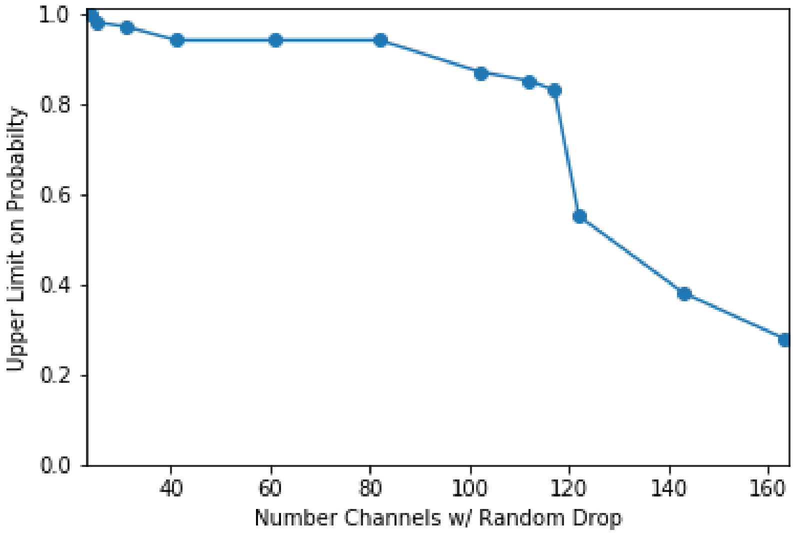

Figure 50. Power injection direction signals: upper bound on the probability of packet drop versus number of channels

Fourth, we evaluated the upper bound on the probability of a packet drop versus the number of bidirectional communication channels between the local controllers and their associated DERs with random packet drop (Figure 51). This category of communication channels is much less resistant to the impacts of packet drop than the other types. Only when all bidirectional channels have at most 0.15 probability of packet drop is the DERMS still functional. Also, its threshold of a sudden drop in the upper bound on the probability occurs at only 25 (15%) channels. However, the communication infrastructure between a local controller and its associated DER(s) is much complicated than the other categories, due to the proximity of the local controller to the DER(s).

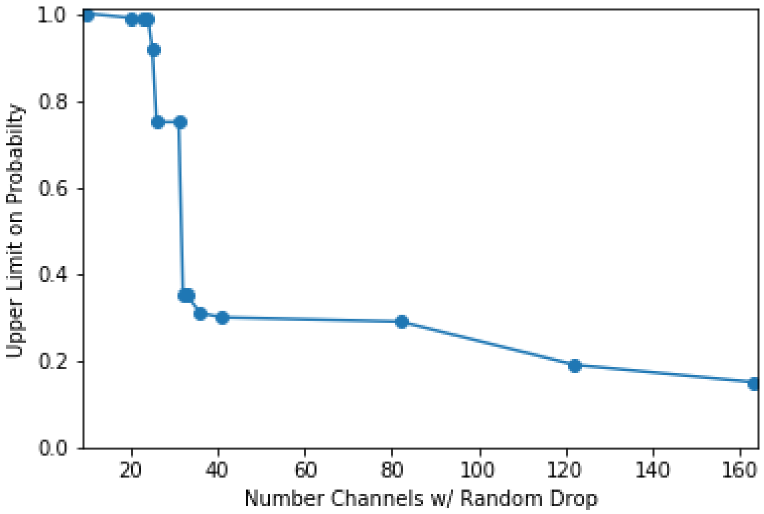

Figure 51. Local controllers and DERs: upper bound on the probability of packet drop versus number of bidirectional channels

2.2.2.6 S4: Random Communication Link Failure

In communication networks, links can randomly fail for several minutes at a time. We simulated this phenomenon similarly to the random packet drop simulations. If a communication link is functional, after it is used to send a message, a Bernoulli random variable is used to decide whether the link will become nonfunctional. If it becomes nonfunctional, then the communication channel cannot be used to send messages for a set amount of time. We evaluated each channel category separately and assumed all other communications channels are always functional. For each channel category we evaluated the upper bound on the probability as a function of the duration that the link will be nonfunctional, assuming every link in the channel category can become nonfunctional.

Within HELICS, the random link failures were simulated within their specific agents: the grid service measurement link failures were simulated within the feeder agent; the power injection directions signal link failures were simulated within the DERMS coordinator agent; the local controller and DER link failures were simulated within the local controller agents. The time step granularities were the same as in the random packet drop simulations.

Figure 52 and Figure 53 give the upper bound on the probability versus the duration of link failure for the feeder head power and voltage magnitude measurements, respectively. Both have a very fast drop from the interval of 1 to 5 minutes, and then level off to a linear-logarithmic decrease. The feeder head power measurements have a steeper decline than the voltage magnitude measurements. This is likely due to each phase of the feeder head powers only having a single point of failure, while each phase of the voltage magnitude measurements have many measurement nodes, each strongly correlated in the same phase.

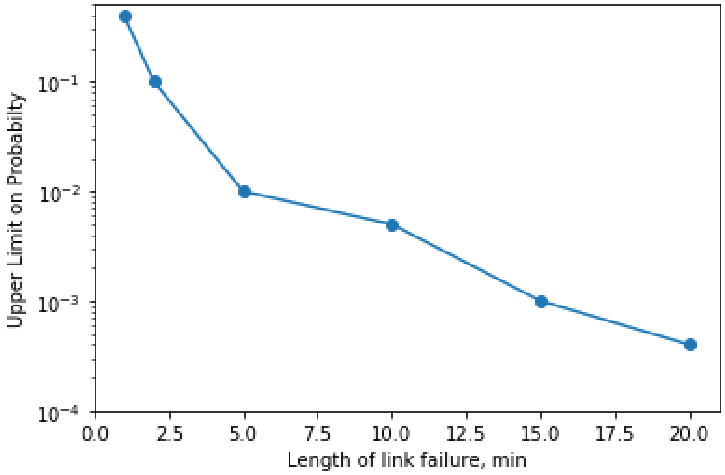

Figure 52. Feeder head power measurements: upper bound of the probability of link failure versus duration of link failure

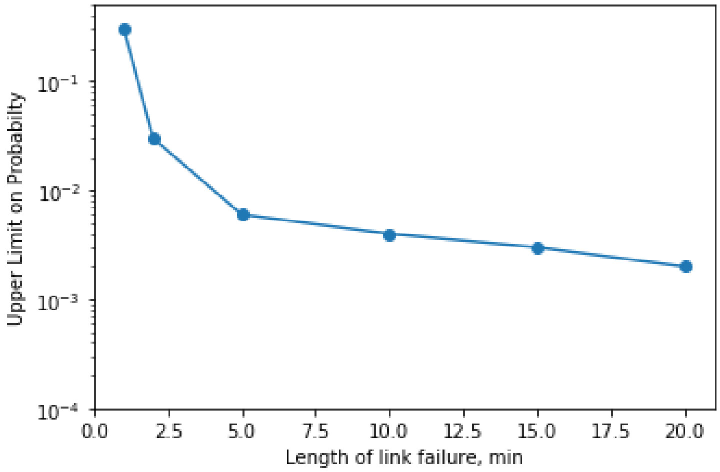

Figure 53. Voltage magnitude measurements: upper bound of the probability of link failure versus duration of link failure

Figure 54 gives the upper bound of the probability versus the duration of link failure for the power injection direction signals. The decrease of the upper bound of the probability being more gradual than all other communication channel categories is likely because each power injection direction signal carries a mix of information about the VPP and voltage magnitude bounds; thus, when some communication links are down, others still carry similar information.

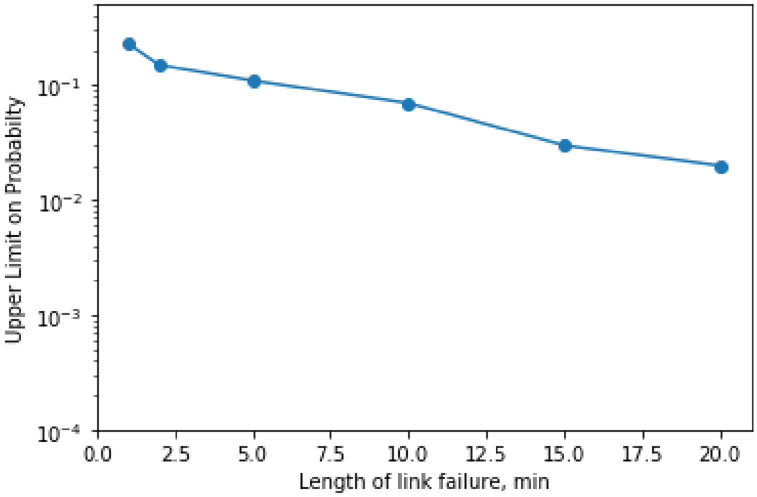

Figure 54. Power injection direction signals: upper bound of the probability of link failure versus duration of link failure

Figure 55 gives the upper bound of the probability versus the duration of link failure between the local controllers and DERs. The gradual leveling off is likely due to the DERs getting their set points stuck when a link fails, which was probably near where it would be if the link was operational.

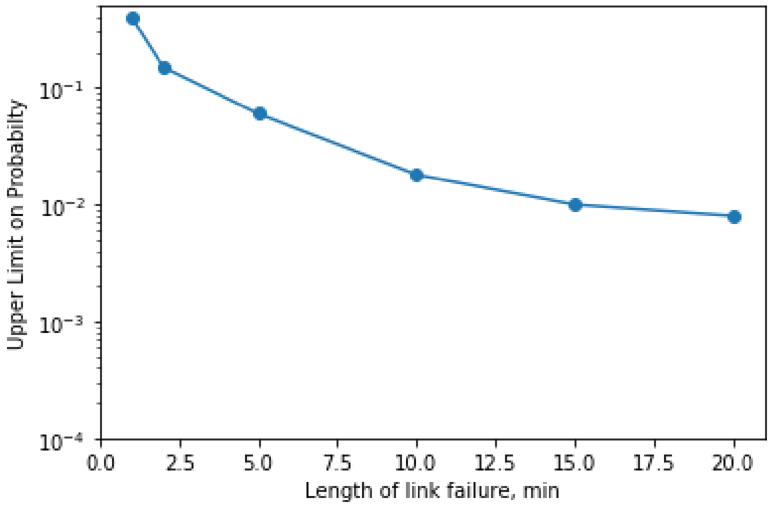

Figure 55. Local controllers and DERs: upper bound of the probability of link failure versus duration of link failure

2.2.2.7 S5: Communication Delays

When sending messages across a communication network, there is an interval of time from when the data in the message was created to when it is used in the next control component, called a delay. This can have adverse oscillatory effects on controlling systems due to the feedback from delayed actions. We evaluated each channel category separately and assumed all other communications channels had no delays. For each channel category, we evaluated the upper bound of the delay that can keep the DERMS functional by applying the delay to all channels within the category.

Within HELICS, the communication delays for the grid service measurements are simulated by adding a native filter to the sending endpoint of the communication channels from the feeder to the DERMS coordinator. The communication delays for the power injection direction signals, sending PQ set points, and receiving PQ measurements, were simulated within the local controller agent. The time step granularity of the grid service measurement communication and power injection direction signal communication was set to 10 seconds, and communication between the local controllers and the DERs was set to 2 seconds so the delays could be evaluated in 10s and 1s of seconds, respectively, instead of minutes and 10s of seconds, as would have been the case if the granularity was set to 1 minute and 10 seconds, respectively.

The upper bounds on the delays for each category are given in Table 12.

Table 12. Upper Bounds on Delays

Communication Channel Category	Upper Bound on the Delay (seconds)
Feeder head power measurements from the substation to the DERMS coordinator	80
Voltage magnitude measurements from various nodes on the feeder to the DERMS coordinator	60
Power injection direction signals from the DERMS coordinator to the DERMS local controllers	50
Power injection set point from DERMS local controllers to DERs	18*
Power injection measurements from DERs to DERMS local controllers	<0.2

*Local controller time step granularity to update set point must be larger than delay.

It is fortunate that the communication channels that are most resistant to the effects of delays are the ones where delays are more common because the origin and destination points must pass through a wide area network.

The upper bound on the communication delay for sending the set points from the local controller to the DER must be less than the time step granularity of the local controller updating the set point; otherwise, the measured power injection used as the reference point in the set point update will be significantly off and cause major oscillations. Thus, the time step granularity of a local controller should be set with this type of delay accounted for, which can also include the delay for the DER to physically implement the received set point.

Only a very small upper bound on the delays could be found for the power injection measurements from DERs to DERMS local controllers. There are two main reasons for the adverse effects that this delay causes (Comden, Wang, and Bernstein 2023b):

1. The local controller uses the received measured power injection as its reference from which to move. A delay causes the local controller to base its decision off an old value that may not reflect the current power injection. The local controller gets stuck in very frequent oscillations.
2. The auto-tuning procedure makes its decisions based off the previous two measured power injections. Oscillations from reason #1 cause the step size to quickly decrease until the local controller has practically no ability to change the set point of the DER.

One possible way to avoid this issue of a power injection delay is to have the power injection measurements sent to the local controller at a higher frequency so that the controller can make sure the previously sent set point is close to the measurement it is receiving and that it is successfully being held there.

3 HIL Demonstration

3.1 HIL Test Setup

The overall schematic of the hardware-in-the-loop (HIL) setup for this project is presented in Figure 56 (a), and the laboratory experiment setup is presented in Figure 56 (b). This setup includes quasi-steady state simulation in OpenDSS, real-time simulation in OPAL-RT, the RT-OPF DERMS algorithm simulated in the PC, hardware inverters under test, and the DERMS local controller host in the NVIDIA DER chip. This HIL platform validates the grid-edge DERM local control algorithm embedded in the NVIDIA DER chip, which will be inserted into the next generation of smart meters. The goal is to evaluate the DER chip's capabilities of dispatching and communicating with real hardware inverters to enable the hardware inverters to collectively achieve the system-level objectives (e.g., voltage regulation and VPP (2030.11, 2021)).

The PHIL portion emulates a single node—including its DERs and associated inverters—which is controlled by the DER chip in a distribution network. The co-simulation portion, primarily the main co-simulation PC, simulates the rest of the distribution network, including the DERMS coordinator and other local controllers. The DERMS coordinator communicates with the DER chip through standard industry protocols (e.g., Distributed Network Protocol [DNP3]) to send power injection directions corresponding to the grid services. Additionally, the DER chip receives measurement data from the inverters (e.g., power). The DER chip calculates the power set points based on the received power injection directions and the measurement data, the local objectives, and the physical bounds, and then it sends the set points to the inverters to be implemented (Wang 2020).

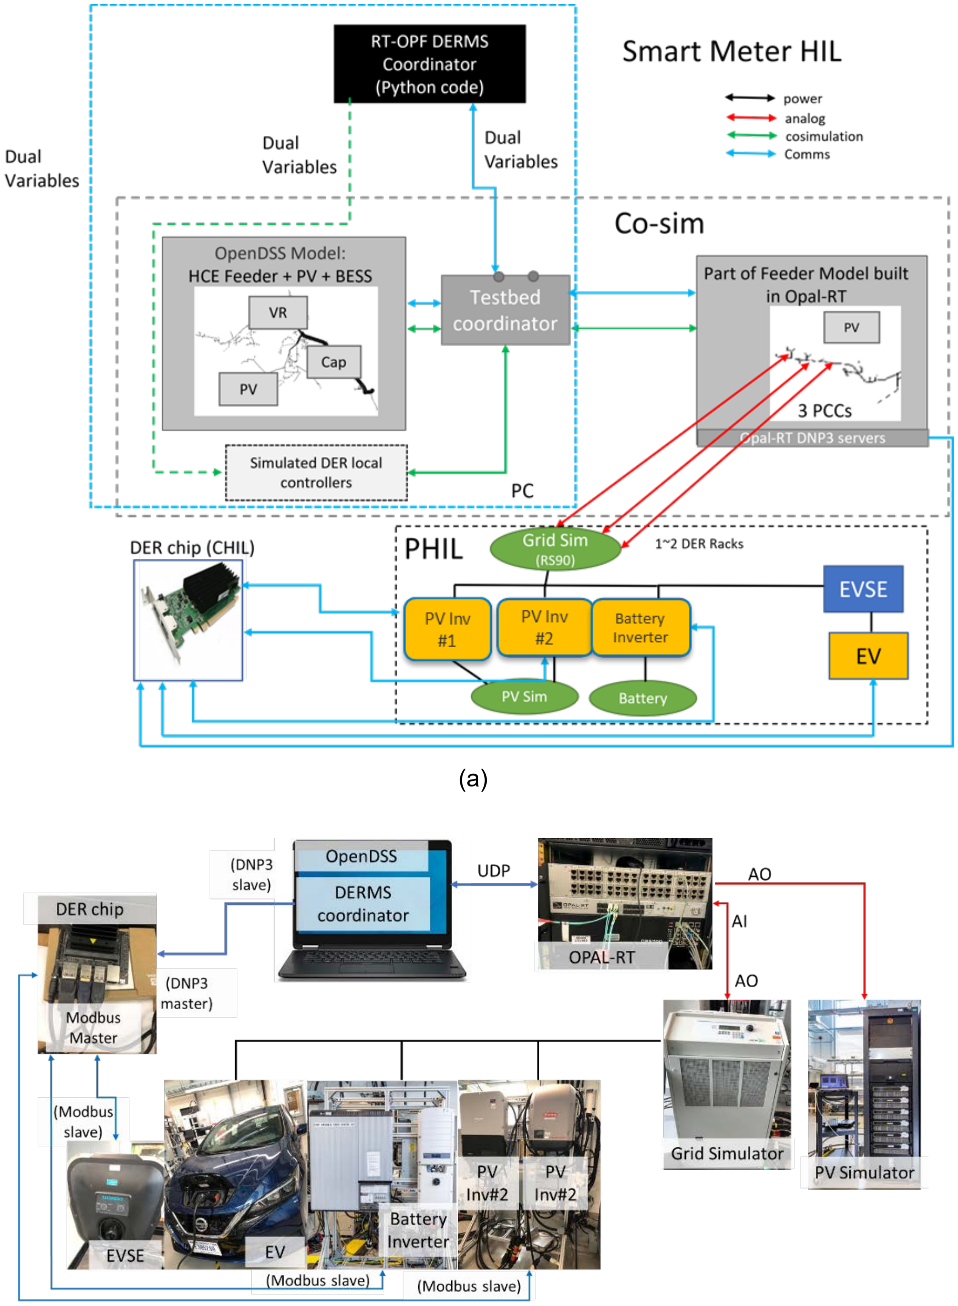

(a)

(b)

Figure 56. (a) Schematics of the HIL setup and (b) laboratory experiment setup

(Comden, Wang, and Bernstein 2023c)

3.1.1 Co-Simulation Platform

The simulation of the distribution feeder and DERMS is primarily run on the main co-simulation PC (see Figure 56) with the use of HELICS (Comden, Wang, and Bernstein 2023c), which controls the global simulation clock and the information passed between the simulation components.

The three-phase distribution feeder and DERs are simulated in OpenDSS every 2 seconds as a quasi-steady-state time series and are interfaced in Python for HELICS with OpenDSSDirect. The feeder model is based on a real-world feeder in Colorado with approximately 2,000 nodes (NVIDIA 2023; Wang et al 2021a). The feeder model sends the grid measurements (nodal voltage magnitudes and feeder head power) to the DERMS coordinator every 1 minute, which is the same frequency at which the coordinator sends the power injection direction signals to the local controllers. The local controllers receive the power measurements from the DERs in OpenDSS every 2 seconds and send the power set points to the DERs every 5 seconds. OPAL-RT is the real-time simulator interfacing with hardware inverters through the grid simulator; it also passes the solar irradiance information from OpenDSS to the PV emulator (Jing, 2021).

A special agent in HELICS was created to replace the local controller embedded in the DER chip. This agent's specific purpose is to communicate the power injections of the DER inverters (through Opal-RT) to OpenDSS, and communicate the voltage magnitude at the DERs' locations on the feeder simulated in OpenDSS to Opal-RT. The communication between OpenDSS and OPAL-RT is through User Datagram Protocol (GitHub 2023). Figure 57 shows the schematics of the co-simulation platform.

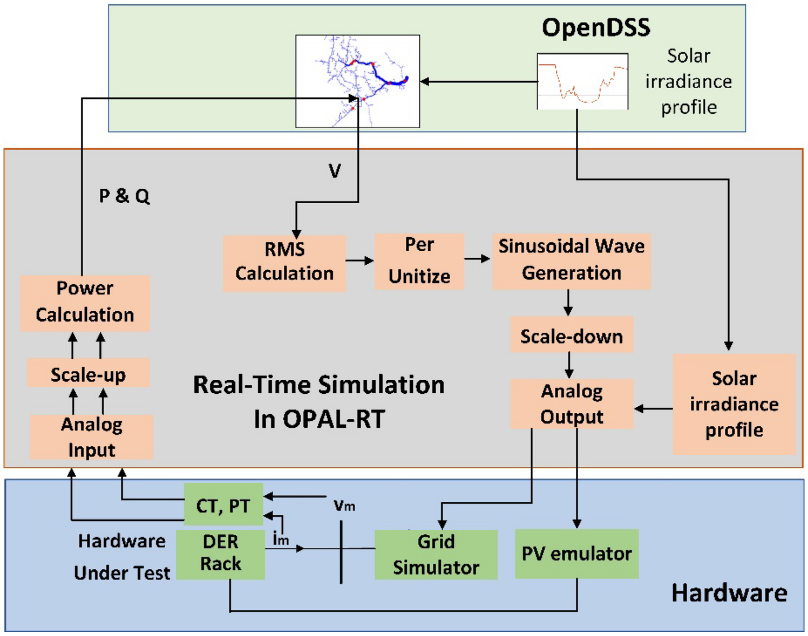

Figure 57. Schematic diagram of the co-simulation between OpenDSS and OPAL-RT

(Comden, Wang, and Bernstein 2023c)

3.1.2 DER Chip

The grid-edge DER chip is a programmed NVIDIA Jetson Nano (NVIDIA 2023), which is a small inexpensive computer that was originally designed for artificial intelligence applications. It is run as it would in a smart meter, where it receives power injection direction signals from the DERMS coordinator and then decides the power set points for the four hardware DERs, including two PV inverters, one battery inverter, and one electric vehicle and electric vehicle supply equipment (EVSE).

The DER chip communicates with inverters from four different manufacturers, which resembles the possible scenario of a resident owning DERs to achieve net-zero energy. For each of the four DER inverters for which the chip decides power set points, it runs a separate instance of the adaptive local controller code (Comden, Wang, and Bernstein 2023c) because each DER has different properties and communication latencies. Additionally, it ensures that their aggregated power has a magnitude less than 12 kW to respect the rating of the point of common coupling (PCC) circuit breaker of the home/lab.

3.1.3 DER Asset and PHIL Interface

PHIL evaluation is used to allow the DER chip dispatch real hardware inverters to output the desired amount of power. As shown in Figure 56, there are three DER hardware inverters and one EV with EVSE. All four DER assets are installed in one DER rack with a common PCC. A high-level description of each DER asset is presented in the following table.

Table 13. DER Properties

Device	Capacity	Dispatch Method
PV Inverter #1	3.8 kVA	External, PQ mode
PV Inverter #2	5 kVA	External, PQ mode
Battery Energy Storage Inverter	5 kVA and 8 kWh	External, P mode
EVSE	4.8 kW and 22 kWh	External, P mode

Figure 57 shows the PHIL interface algorithm implemented in OPAL-RT. For the PHIL setup, the root-mean-square value of the voltage where the four DER assets are virtually connected is sent from OpenDSS to OPAL-RT. This voltage is per unit, and then regenerated as sinusoidal waveform, scaled down, and then sent to the grid simulator (RS45) through analog output. The solar irradiance is in per unit and also sent from OpenDSS to OPAL-RT and then to the PV emulator. The DER chip dispatches all four DER assets, and the lumped current of the four DER assets and the PCC voltage are sent back to OPAL-RT through the current transducer and potential transducer. Inside OPAL-RT, this current and voltage from the hardware inverters are converted to physical values and are used to calculate the physical power at the PCC of the DER rack. Finally, this active and reactive power of the DER rack is sent from OPAL-RT to OpenDSS to drive the controlled source, thus closing the PHIL simulation to reflect the virtual voltage in OpenDSS to the hardware inverters and the actual power from the DER assets to the OpenDSS controlled source (Padullaparti 2022; Wang 2021).

3.1.4 Communication Between Elements

All the communication among the elements in the lab is over Ethernet using a subnet.

- The DERMS coordinator, in the main co-simulation PC, sends the power injection directions signals, PV forecast, and simulation time stamp to the DER chip with DNP3 using OpenDNP3 (GitHub 2023). The main co-simulation PC connects to the DER chip by having the PC run a DNP3 outstation server and the DER chip run a DNP3 master server. The DERMS signal receiver collects the signals received from the coordinator in HELICS and appends the values to a local text file one by one. The outstation server sends the appended values one by one to the master server. The master server appends a text file local to the DER chip as it receives values from the outstation server. The local controller code on the DER chip reads the last 10 lines of the text file every 10 seconds to update its coordinator signal information.

- The test bed coordinator, in the main co-simulation PC, interfaces with the feeder model in HELICS and communicates with Opal-RT via User Datagram Protocol. The feeder model sends the resulting voltage magnitude of the HIL-emulated feeder node and the solar irradiance to OPAL-RT. The voltage magnitude is then used in OPAL-RT to generate the sinusoidal waveform for the grid simulator, and the solar irradiance is sent from OPAL-RT to the PV emulator. Both signals are sent through analog output channels. The four DER hardware components are connected at the same PCC; thus, the PCC voltage and lumped current are sent back to OPAL-RT. The lumped active and reactive power at the PCC is then calculated in OPAL-RT and sent to the OpenDSS simulation model to close the PHIL loop (Wang et al. 2020a).
- The DER chip sends the power set points to the inverters and receives the measurement data via Modbus Transmission Control Protocol/Internet Protocol. This is set up through Modbus (PyPI 2023), and the correct registers must be used to write/read the inverter. The data communicated and the frequency between the chip and each inverter are different because of the type of DER and latency. Table 14 gives the frequency at which the set points were sent to the inverter and the measurements were sent to the DER chip. It also gives the variables sent, including the active power (P), reactive power (Q), and state of charge. Note that, because the register of PV Inverter #1 used to send the reactive power measurement is not functional on our device, we estimate the reactive power measurement with its set point.

Table 14. DER Communication

Device	Frequency (s): Set/Measure	Variable: Set/Measure
PV Inverter #1	15/6	P, Q/P
PV Inverter #2	5/2	P, Q/P, Q
Battery Energy Storage Inverter	10/4	P/P, State of charge
EVSE	10/4	P/P

3.2 HIL Test Scenarios

In total, we ran three use cases, including smooth solar irradiance with changes in VPP references, fluctuated solar irradiance, and the big change in the system voltages caused by the change of the LTC tap position. The results of each use case are described in the next subsection.

3.3 HIL Test Results

3.3.1 HIL Test With Changes in VPP Bounds

This test scenario was borrowed and slightly modified from the scenario we created for the automatic step-size tuning where the solar irradiance profile was smooth and the VPP had two different step changes. This scenario is made to showcase the speed at which the DERMS can coordinate its DERs to get the feeder head power in each phase to be within the changed VPP bounds. With regards to the HIL test, we wanted to see how well the DER inverters could follow the set point given by their local controller when the VPP bounds for the phase were changed suddenly. The HIL node is located on the feeder on Phase B. Figure 58 gives the VPP bounds (“Psub_set_lower”, “Psub_set_upper”) and feeder head power (“Psub_meas”) for Phase B over the 2-hour test.

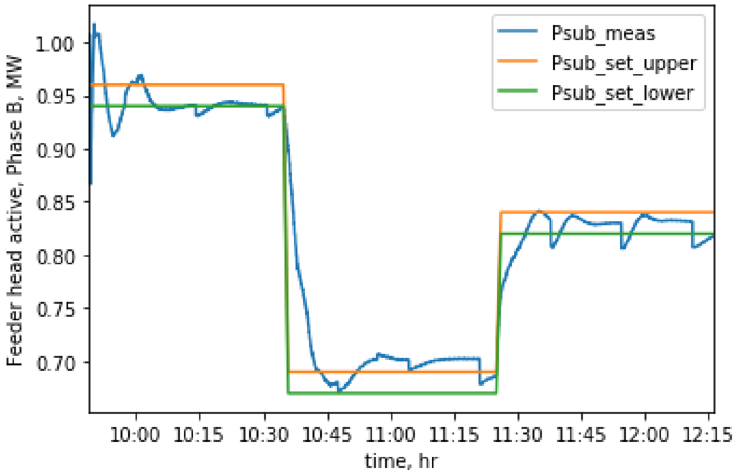

Figure 58. VPP bounds and feeder head power of Phase B

The DERMS coordinator reacts to the changes in the VPP bounds by sending out dual variables associated with the VPP bound and voltage bound violations once every minute to all the nodes. The dual variables indicate the opposite direction that power injections should go. The dual variables sent to the HIL node are shown in Figure 59, separated for active and reactive power. Summing the dual variables for each active/reactive power (“Total”) gives the signal from the coordinator to the local controller that the active power injection needs to be increased after 10:32 and decreased after 11:25 due to the changing VPP bounds, while the reactive power injection needs to be first decreased after 10:32 to counteract the voltage increase from increasing the active power injection and vice versa after 11:25.

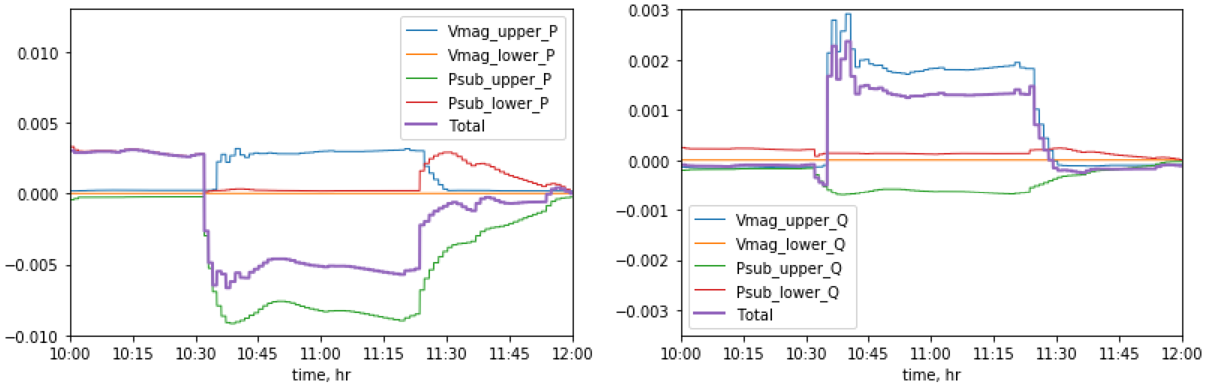

Figure 59. Dual variables calculated by the DERMS coordinator for (left) active power and (right) reactive power

First, we look at PV Inverter #1 controlling the power injection of solar PV. The local controller measures its power and sends a power set point every 6 seconds and updates the set point every 15 seconds. Figure 60 the active power measurements (“Ppv_mes_1”), set points (“Ppv_set_1”), and the forecasted maximum PV power (“PV_Pmax_forecast_1”). The active power injection from 10–10:32 shows a small amount of PV curtailment, which disappears after the VPP bounds are changed at 10:32. The reason for the step pattern in the power measurements is that the set points that are sent over Modbus Transmission Control Protocol must be rounded to the nearest integer percentage of the inverter rating. To get a better look around the moment when the VPP changes at 10:32, we zoom in (Figure 61; Figure 60) to see that the curtailment disappears due to the sent dual variables with respect to the active power.

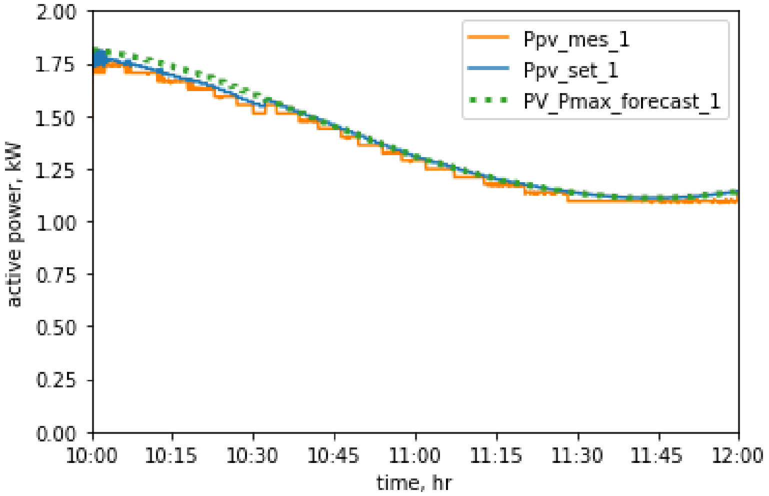

Figure 60. Active power injection of PV Inverter #1

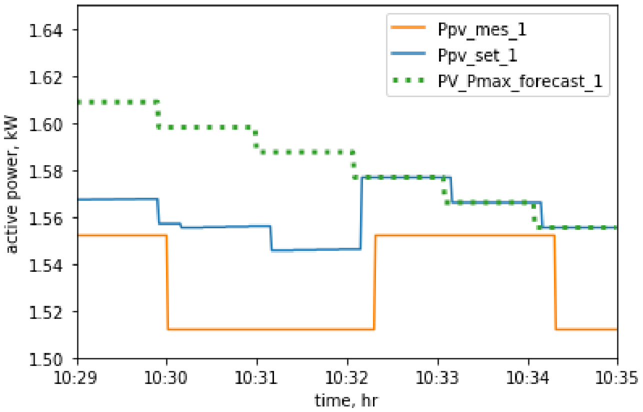

Figure 61. Zoomed-in active power injection of PV Inverter #1 over 6 minutes

The set point of the reactive power from PV Inverter #1 smoothly follows the dual variables, as shown in Figure 62. Unfortunately, the register on the inverter that was supposed to send the reactive power measurement to the local controller was not functional, but we could confirm by observing the display screen of the inverter that the reactive power was relatively close to the set point being sent to it.

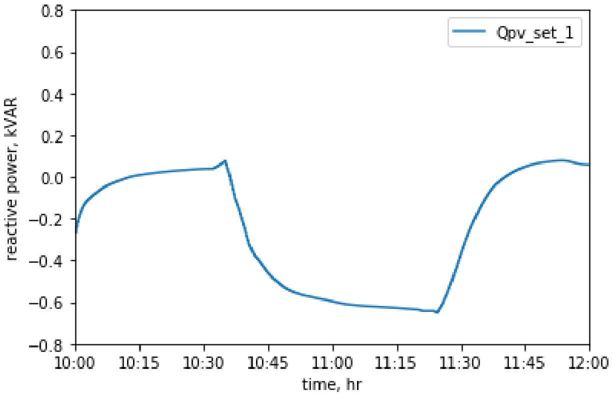

Figure 62. Reactive power injection of PV Inverter #1

Next, we look at PV Inverter #2 controlling the power injection of solar PV. The local controller measures its power and sends a power set point every 2 seconds and updates the set point every 5 seconds. Figure 63 gives the active power measurements (“Ppv_mes_2”), set points (“Ppv_set_2”), and the forecasted maximum PV power (“PV_Pmax_forecast_2”). The PV curtailment pattern is very similar to that of PV Inverter #1 where there is some small curtailment until 10:32 when the VPP bounds change, which removes the curtailment to increase generation due to the changes in the dual variables. The noisy power measurements are typical for these inverters; fortunately, it only minimally impacts the set points decided by the local controller. Figure 64 zooms into the time right before and right after the changes in the VPP. The set point is still able to accurately follow the dual variable signals from the coordinator, even if the actual injected active power has around 7% volatility. Even if the actual active power injection is not fully accurate, Figure 51 shows that on average, it does follow the set point.

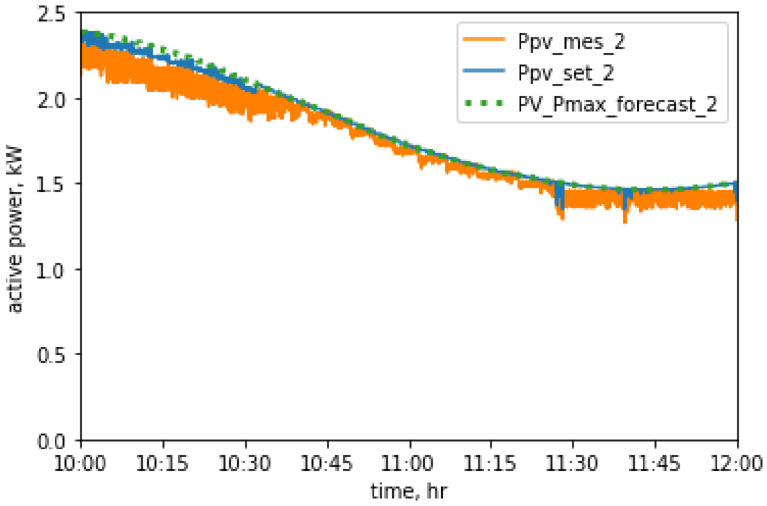

Figure 63. Active power injection of PV Inverter #2

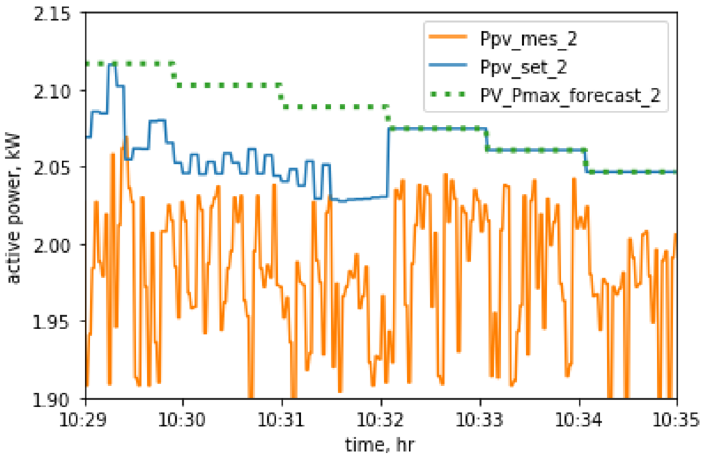

Figure 64. Zoomed-in active power injection of PV Inverter #2 over 6 minutes

Interestingly, the measured reactive power ("Qpv_mes_2") as shown in Figure 65 follows the set point ("Qpv_set_2") closely with much less volatility as compared to the active power. It also smoothly follows the dual variable signals from the coordinator to decrease the voltage by decreasing the reactive power injection at 10:32.

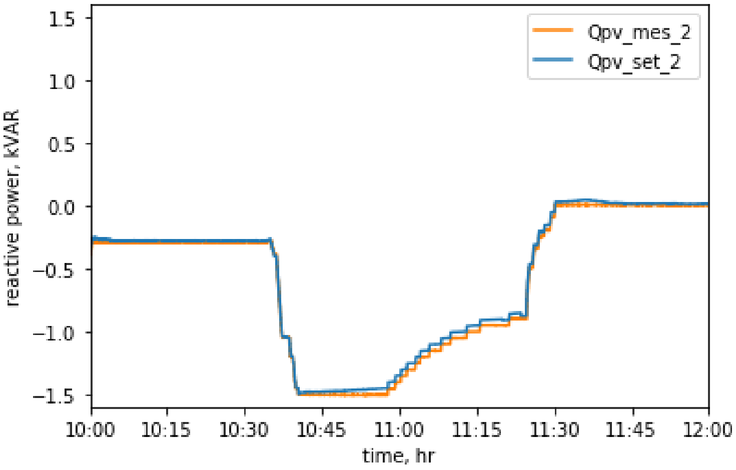

Figure 65. Reactive power injection of PV Inverter #2

The battery energy storage inverter controls the charging and discharging active power of a battery. The local controller measures its power and state of charge and sends a power set point every 4 seconds, then updates the set point every 10 seconds. Because the battery is not concerned with curtailment—as the PVs are—it has more flexibility to adjust with the VPP bound changes. Figure 66 gives the active power measurements ("Pstor_mes") and set points ("Pstor_set"). The power measurements follow the set point very closely. The local controller uses the dual variables as a signal to charge the battery, and then discharge battery when the VPP bounds change, in almost a step change manner as the change in VPP bounds. However, at 10:53, the battery hits its minimum allowable state of charge and can no longer provide energy toward the VPP, as shown in Figure 67.

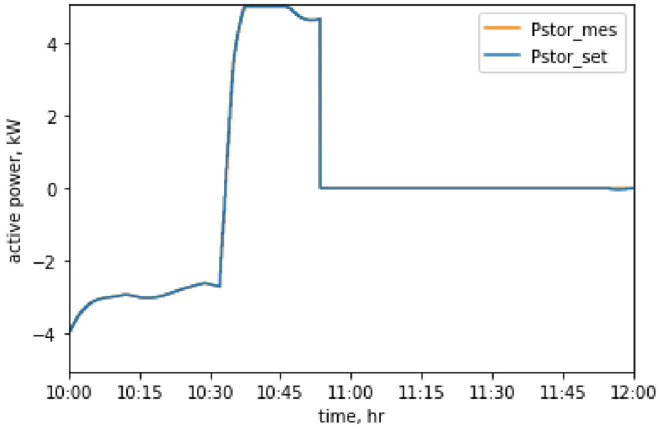

Figure 66. Active power injection of the battery energy storage inverter

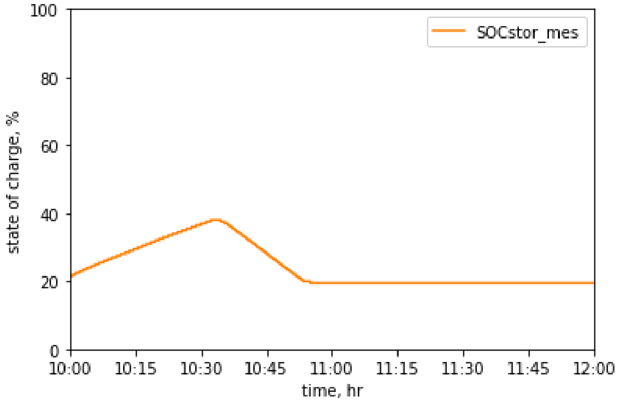

Figure 67. State of charge of the battery controlled by the battery energy storage inverter

The EVSE was used to charge a Nissan Leaf electric vehicle. The local controller measures its power and sends a power set point every 4 seconds, then updates the set point every 10 seconds. The EV has a lower bound on charging of 1.68kW (7 amps at 240 V). Figure 68 gives the active power measurements ("Pstor_mes_evse") and set points ("Pstor_set_evse"). The power measurements follow the set point very closely. Like the battery associated with the Solar Edge inverter, it starts out charging the electric vehicle above the minimum amount until 10:32 when the dual variables make it lower the load and charge at its minimum value of 1.68 kW.

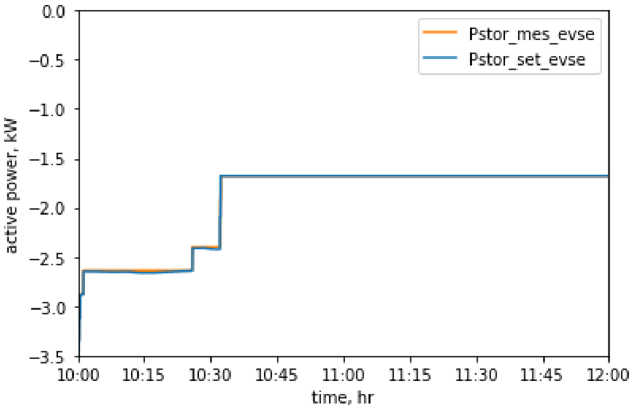

Figure 68. Active power injection of the EVSE

3.3.2 HIL Test With Fluctuating Solar PV Power

This 2-hour test scenario was borrowed and slightly modified from the scenario we created for the automatic step-size tuning where the PV power trajectory starts out smooth and then suddenly becomes volatile at 10:50. This challenges the DERMS to keep the feeder head powers within the VPP bounds. The HIL node is located on the feeder on Phase B.

Figure 69 gives the feeder's total PV power (available and measured after curtailment) and the feeder head power with VPP bounds for Phase B. These sudden PV fluctuations cause the DERMS coordinator to adjust its dual variables signals sent to the HIL node (Figure 70). The dual variables for the active power are primarily used to counteract the PV power fluctuations to keep the feeder head power within the VPP bounds, whereas the dual variables for the reactive power are primarily used to keep the voltage magnitudes within its bounds.

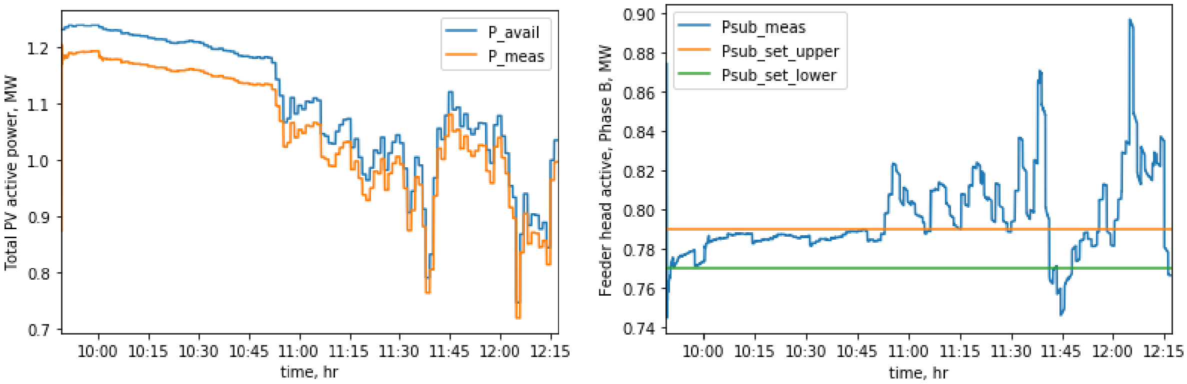

Figure 69. (Left) Total PV active power injected into the feeder and (right) VPP bounds and feeder head power of Phase B

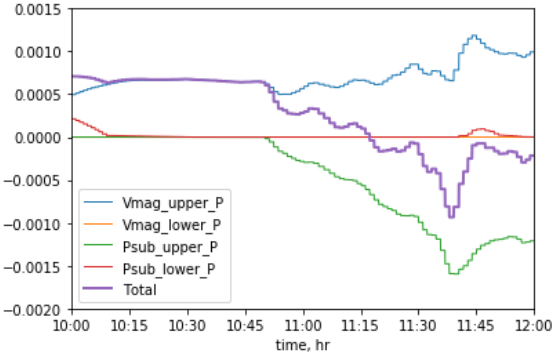

Figure 70. Dual variables calculated by the DERMS coordinator for (left) active power and (right) reactive power

PV Inverter #1, controlling the power injection of the solar PV, mostly uses active power control to smooth out the PV fluctuations and uses reactive power control to counteract the PV fluctuation effects on the feeder voltages (Figure 71).

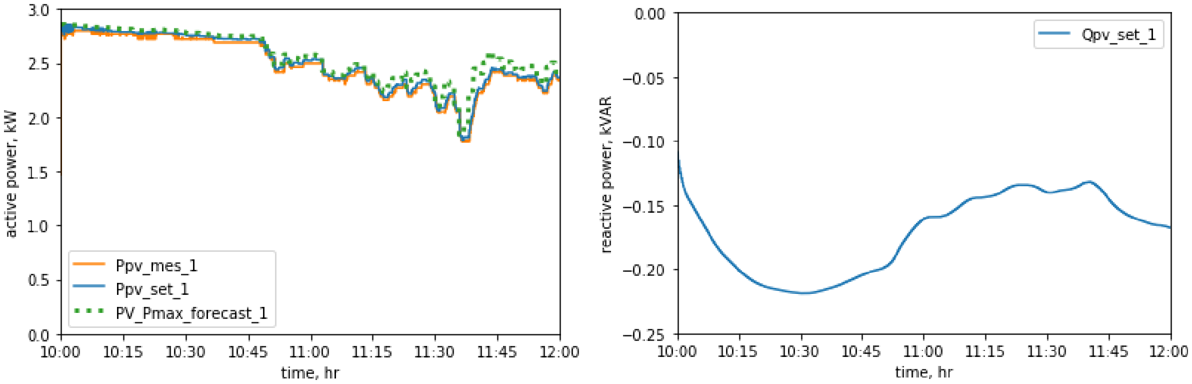

Figure 71. (Left) Active and (right) reactive power injection of PV Inverter #1

PV Inverter #2, controlling the power injection of the solar PV, mostly uses active power control to lessen PV curtailment when PV fluctuates and uses reactive power control to counteract the PV fluctuation effects on the feeder voltages from 11:38–11:45 (Figure 72). Although the reactive power injection sometimes has a bias from its set point, it does follow changes in the set point almost instantly.

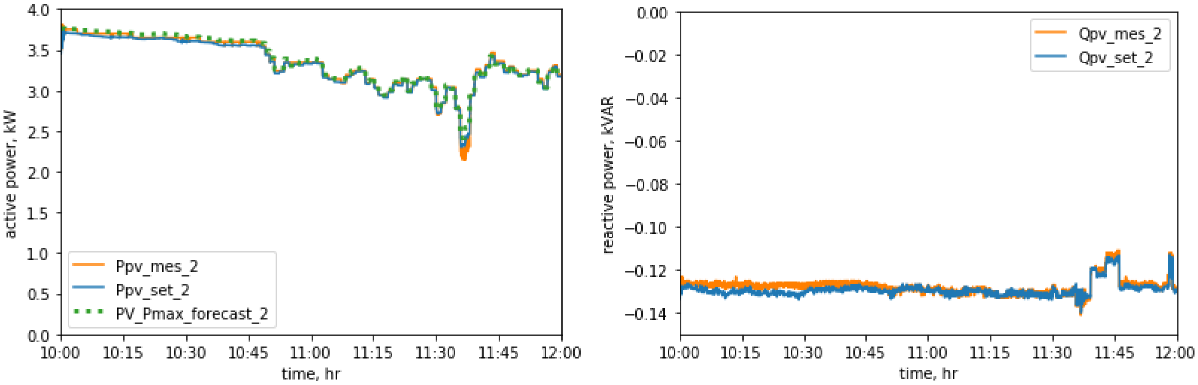

Figure 72. (Left) Active and (right) reactive power injection of PV Inverter #2

It is apparent that the trajectory of PV Inverter #1 is smoother than that of PV Inverter #2. Figure 73 gives a zoomed-in look from 11:37:30 to 11:39:30 of the active power injections; PV Inverter #1 has a longer delay in acting upon the set points than PV Inverter #2. Additionally, PV Inverter #1 has a 10-times longer delay in measuring and sending its power measurements with Modbus Transmission Control Protocol than PV Inverter #2. Figure 74 shows histograms of the combined measurement and communication delays. The combination of a slower measurement read time and slower implementation of set points makes PV Inverter #1 give smoother power injections.

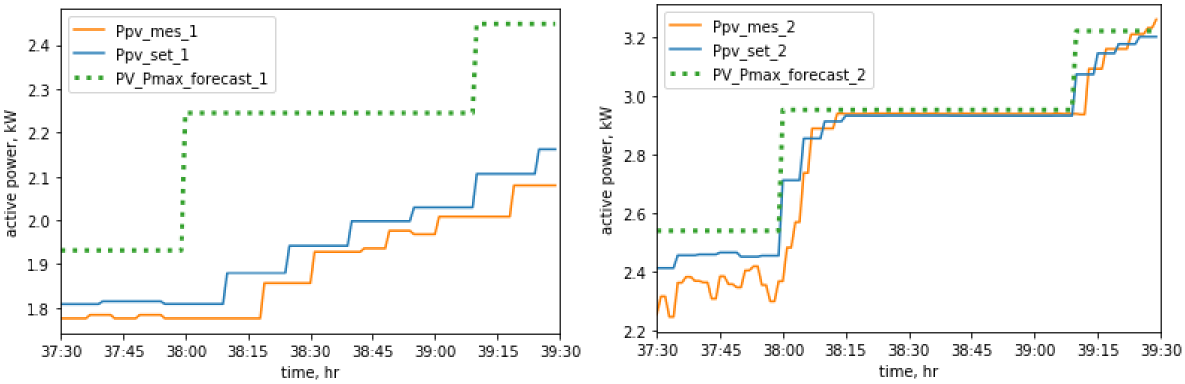

Figure 73. 2-minute zoomed-in view of the active power injections of (left) PV Inverter #1 and (right) PV Inverter #2

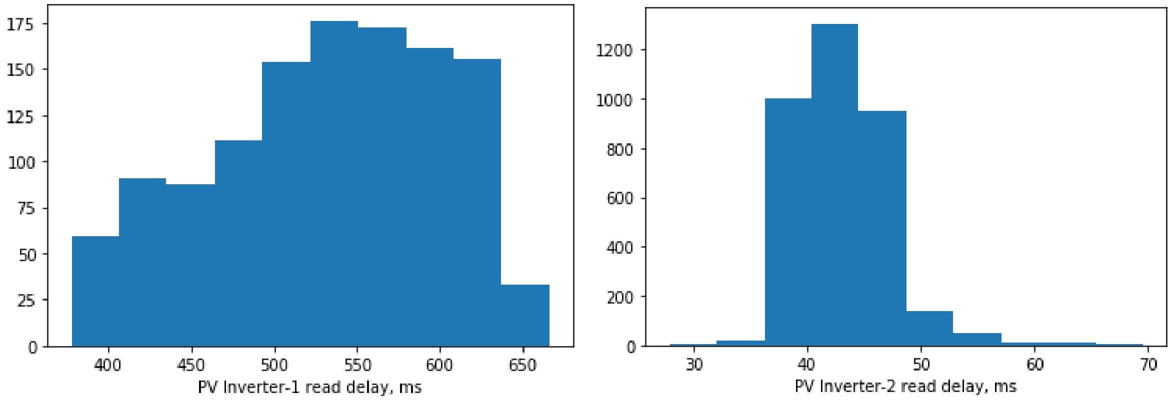

Figure 74. Histograms of the combined measurement and communication delays of (left) PV Inverter #1 and (right) PV Inverter #2

The battery energy storage inverter, controlling a battery, uses its active power injection to counteract the PV fluctuations, even switching from charging to discharging (see Figure 75). However, because the Nissan Leaf was already charging at its minimum charging power before the PV fluctuations happened, it could not lessen charging power to counteract the PV fluctuations and stayed constant during the whole test (Figure 76).

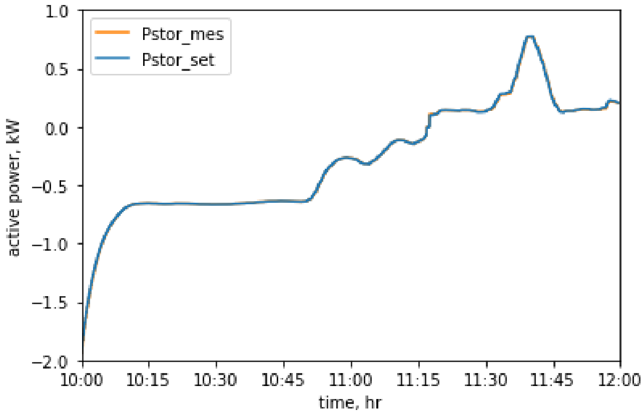

Figure 75. Active power injection of the battery energy storage inverter

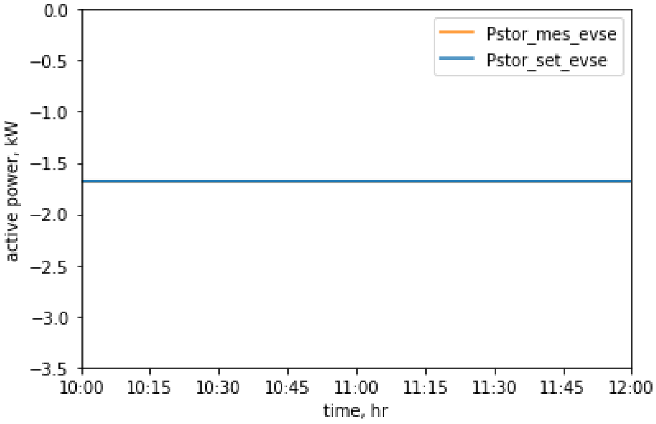

Figure 76. Active power injection of the EVSE

3.3.3 HIL Test With Changing LTC Position

This 2-hour test scenario was borrowed and slightly modified from the scenario we created for the automatic step-size tuning in which the LTC tap positions have large step changes. This results in large step changes in the voltage magnitudes, where the voltage magnitude bound that becomes active changes (Figure 77). The DERMS reacts to the changes by also making step changes in the reactive power dual variables by switching whether "Vmag_lower_Q" or "Vmag_upper_Q" makes up most of the total dual variable magnitude (Figure 78). Simultaneously, the active power dual variables are used to maintain the VPP; for example, during 11:33–12:00, the voltage support and VPP active power dual variables counter each other so that the total active power dual variable is near zero (Figure 78).

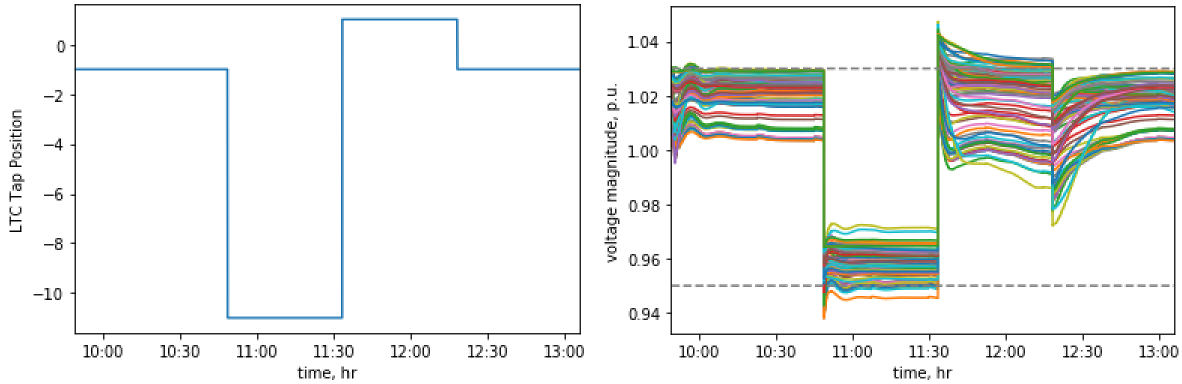

Figure 77. (Left) LTC tap position at the feeder head and (right) voltage magnitudes across the feeder

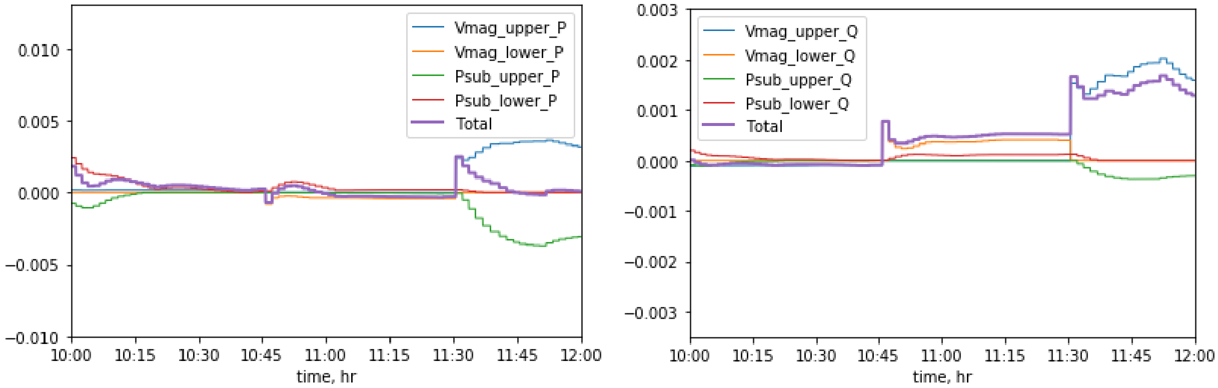

Figure 78. Dual variables calculated by the DERMS coordinator for (left) active power and (right) reactive power

PV Inverters #1 and #2, controlling the power injection of a solar PV, mostly uses reactive power control to contribute to the voltage support (Figure 79 and Figure 80).

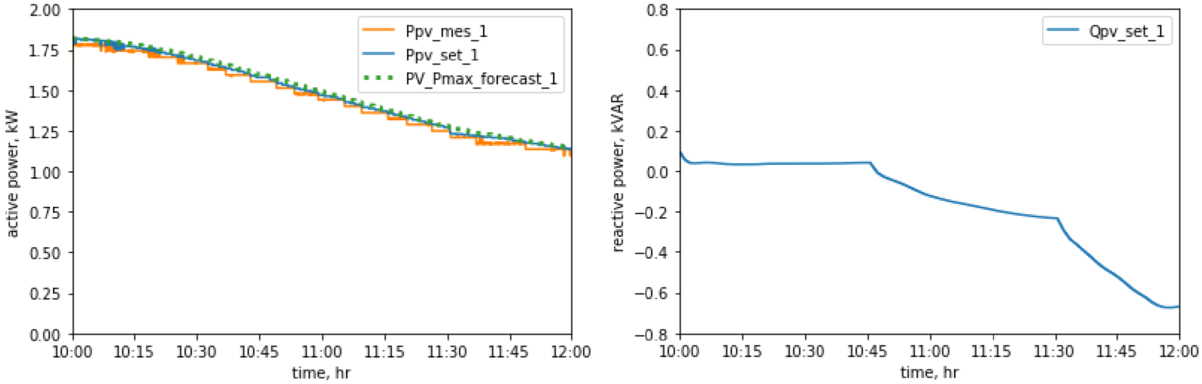

Figure 79. (Left) Active and (right) reactive power injection of PV Inverter #1

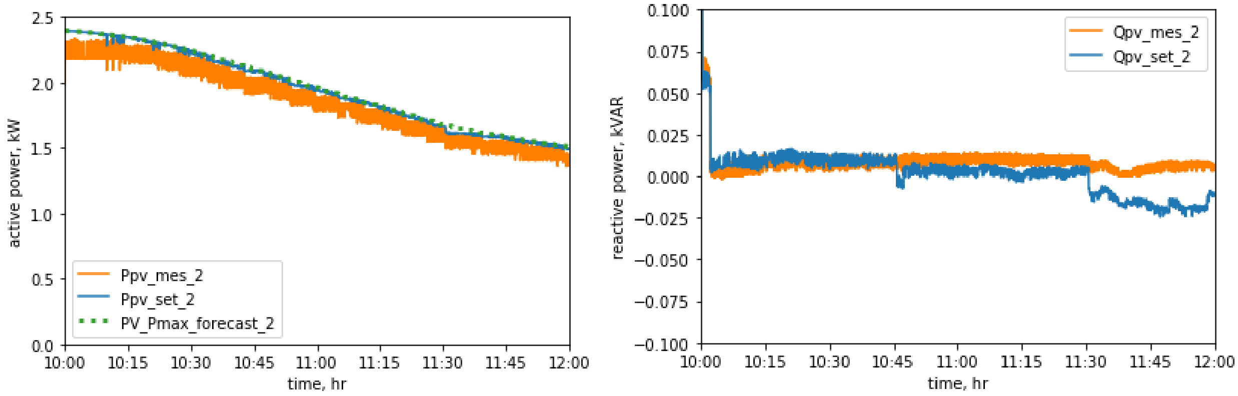

Figure 80. (Left) Active and (right) reactive power injection of PV Inverter #2

The battery energy storage inverter, controlling a battery, uses its active power injection to first contribute to the voltage support and then ease back into maintaining the VPP; this is especially noticeable in Figure 81 during 11:33–12:00 with its initial drop in power before slowly coming back. However, because the Nissan Leaf is controlled with discrete changes in active power, the dual variable signals were not strong enough to change its set point, even though the DER chip did try at 11:33 (Figure 82).

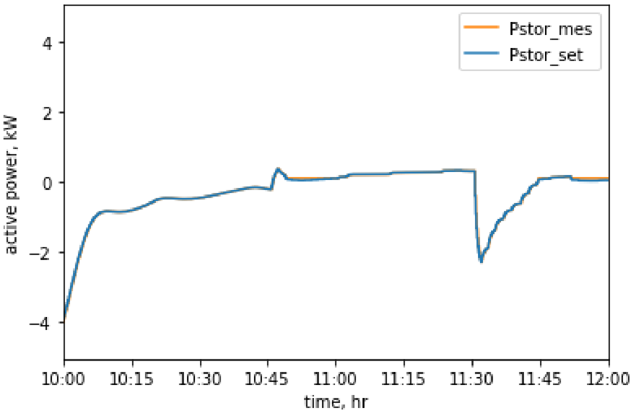

Figure 81. Active power injection of the battery energy storage inverter

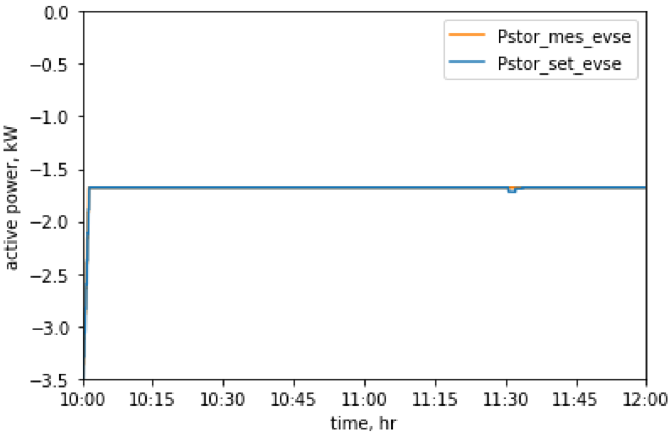

Figure 82. Active power injection of the EVSE

3.4 Summary of HIL Validation

The testing demonstrated the communication and functionalities of the NVIDIA DER chip. To be more specific, these HIL tests demonstrate that:

- Real-world inverters controlling PV and batteries are capable and effective at following set points from an RT-OPF DERMS that automatically tunes itself to local conditions.
- The local controller code is lightweight enough to be placed on a Jetson Nano micro-computer and effectively control four separate inverters simultaneously.
- The DERs can be controlled to quickly react when the VPP set points change and when renewable energy power injections anywhere in the feeder become volatile.

4 Conclusion

NREL's smart-meter-based RT-OPF DERMS provides a unique real-time, distributed, and plug-and-play optimization platform that coordinates the operation of massive numbers of DERs to ensure voltage and power quality, maximize social welfare, and emulate VPPs. This project was able to push RT-OPF DERMS closer to commercialization in the following three ways:

1. The control algorithm underlying RT-OPF DERMS was enhanced with automatic tuning so that real-world implementation and function requires much less labor. Based on our previous project experience, tuning the RT-OPF usually takes at least 2 months. In addition, the adaptive tuning makes the RT-OPF still functions well under changing grid conditions. Therefore, the adaptive automatic tuning removes the largest barrier for commercializing RT-OPF DERMS.
2. Upper bounds on communication issues (packet losses, link failures, and latency) were investigated so that the communication requirements for RT-OPF DERMS can be outlined. With that a least-cost communication system can be chosen based on the performance requirements of the RT-OPF DERMS. This will help DERMS deployer a significant amount of budget as communication infrastructure is always very costly. Thus, this removes the second largest barrier for commercialization.
3. The functionality and communication capability of the DER chip is demonstrated successfully through HIL evaluation with hardware inverters and real EV. This creates the conditions for the next step, which is to insert the DER chip into the smart meter hardware and develop the final product. Therefore, the concept of a smart-meter-based DERMS was shown to be a viable option for commercialization.

In addition, with the contributions from the two industry partners, Utilidata and NVIDIA, and our advisor board, multiple inverter manufactures (CE+T, Enphase, SMA and SolarEdge), communication architecture of RT-OPF DERMS is also defined. This helps the investors (e.g., control vendor and utilities) better understand the potential investment in communication infrastructure. In this TCF project, the technology maturity improvements are successfully carried out to move the technology forward for product development and market, however, our industry partner didn't continue the commercialization effort due to the change of their company's mission and strategy. Therefore, we continue to initiate collaborations with other companies.

5 Path Forward

Utilities are commercially motivated to invest in infrastructure, and regulators are trying to ensure that utilities make prudent investments that prepare them for a dynamic, DER-heavy grid. While the project pushed NREL’s smart-meter-based RT-OPF DERMS closer to commercialization, the following are some opportunities for further study to address some key challenges that the industry faces:

- Develop the more advanced communication protocols: the communication protocol IEEE 2030.5 is commonly used between grid-edge controller and DER inverters. To make the DERMS have more market, it is necessary to develop IEEE 2030.5 in the local controller as well.
- End-user interaction with the DERMS: RT-OPF DERMS requires that the owner of each DER put their preferences via settings into its associated local controller. Further market study must be done to design preference alignment and incentive programs that will attract DER owners to participate based on their goals. Also, if end users sign up to a demand response program that exists outside of the DERMS, it will be important to study how this would affect the performance of RT-OPF DERMS.
- Demonstrate benefits provided by DERMS: this can be achieved by using real-world grid service signals and tariff information. The possible grid service includes voltage regulation, peak loading shaving, VPP, energy arbitrary, and fast frequency response. And testing should be covering representative scenarios of a year, and a techno-economic analysis can be performed to illustrate the dollar values from the DERMS. This gives straightforward information for decision-making.
- Respond to FERC Order 2222: This order sets requirements for wholesale markets and requires that DERs should aggregate together with more than 100 kW to participate the market. The DERMS should response to this order and embody the aggregated DER participation. This way DERMS can be treated as a third-party DER aggregator to coordinate with utility ADMS and DERMS system for grid services. Investigation shows that this will be a feasible pathway for utilities to efficiently integrate DERs’ contributions and manage them for a win-win situation.

To move forward, the potential commercialization partners will be grid-edge and DERMS vendor, especially smart meter companies. Their platforms can be used to host the RT-OPF DERMS algorithm and platforms’ SCADA, communication capabilities, and other resources can be leveraged for efficient commercialization and product development.

References

Bernstein, A., and E. Dall'Anese. 2017. "Linear Power-Flow Models in Multiphase Distribution Networks." Presented at the IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT-Europe). Torino, Italy.

Bernstein, A., et al. 2020. "Real-Time Optimization and Control of Next-Generation Distributed Infrastructure." Presented at the Network Optimized Distributed Energy Systems (NODES) Workshop and Demonstration. January 16, 2020. https://www.nrel.gov/docs/fy21osti/78742.pdf.

Comden, J., J. Wang, and A. Bernstein. 2022. Software Record SWR-22-50: Automatic Tuner for the Step Sizes of Gradient-Based RT-OPF DERMS Control.

Comden, J., J. Wang, and A. Bernstein. 2023a. "Adaptive Primal–Dual Control for Distributed Energy Resource Management." *Applied Energy* 351: 121883.

Comden, J., J. Wang, and A. Bernstein. 2023b. "Evaluation of Communication Issues in Primal-Dual-Based Distributed Energy Resource Management Systems (DERMS)." Presented at the 2023 IEEE Power & Energy Society General Meeting (PESGM). Orlando, Florida, July 16–20.

Comden, J., J. Wang, S. Ganguly, S. Forysth, R. Gomez, and A. Bernstein. 2023c. "Hardware-in-the-Loop Evaluation of Grid-Edge DER Chip Integration into Next Generation Smart Meter." Presented at the 2023 IEEE International Conference on Communications, Controls and Computing Technologies for Smart Grids (SmartGridComm). Glasgow, Scotland, October 31–November 3, 2023.

Comden, J., J. Wang, and A. Bernstein. 2023d. "Study of Communication Boundaries of Primal-Dual-Based Distributed Energy Resource Management Systems (DERMS)." Presented at the 2023 IEEE Power & Energy Society Innovative Smart Grid Technologies (ISGT).

Dall'Anese, E., S. S. Guggilam, and A. Simonetto. 2017. "Optimal Regulation of Virtual Power Plants." *IEEE Transactions on Power Systems* 33 (2): 1868–81.

GitHub. 2023. "OpenDNP3." https://github.com/dnp3/opendnp3.

Goldstein, T., M. Li, and X. Yuan. 2015. "Adaptive Primal-Dual Splitting Methods for Statistical Learning and Image Processing." *Advances in Neural Information Processing Systems* 28.

Institute of Electrical and Electronics Engineers (IEEE). 2021. IEEE Std 2030.11TM-2021 – IEEE Guide for Distributed Energy Resourcese Mangament Systems (DERMS) Functional Specification.

NVIDIA. 2023. "Jetson Nano Developer Kit." https://developer.nvidia.com/embedded/jetson-nano-developer-kit.

Obi, M., T. Slay, and R. Bass. 2020. "Distributed Energy Resource Aggregation Using Customer-Owned Equipment: A Review of Literature and Standards." *Energy Reports* 6: 2358–69.

Padullaparti, H., et al. 2022. “Evaluation of Data-Enhanced Hierarchical Control for Distribution Feeders with High PV Penetration.” *IEEE Access* 10: 42860–72.

PyPI. 2023. “Modbus Test Kit.” https://pypi.org/project/modbus-tk/.

Seal, B., A. Renjit, and B. Deaver. 2018. *Understanding DERMS*. Palo Alto, CA: Electric Power Research Institute. https://www.epri.com/#/pages/product/000000003002013049/?lang=en.

Wang, J. 2021. “Voltage Regulation Performance Evaluation of Distributed Energy Resource Management via Advanced Hardware-in-the-Loop Simulation.” *Energies* 14 (20).

Wang, J., et al. 2020a. “Performance Evaluation of Data-Enhanced Hierarchical Control for Grid Operations.” Presented at the 2020 IEEE Power & Energy Society General Meeting (PESGM).

Wang, J., et al. 2020b. “Performance Evaluation of Distributed Energy Resource Management via Advanced Hardware-in-the-Loop Simulation.” Presented at the 2020 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT).

Wang, J., et al. 2021a. “Hardware-in-the-Loop Evaluation of an Advanced Distributed Energy Resource Management Algorithm.” Presented at the 2021 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT).

Wang, J., et al. 2021b. “Performance Evaluation of an Advanced Distributed Energy Resource Management Algorithm.” Presented at the 2021 IEEE International Conference on Communications, Controls and Computing Technologies for Smart Grids (SmartGridComm).

Wang, J., J. Huang, and X. Zhou. 2020. “Performance Evaluation of Distributed Energy Resource Management Algorithm in Large Distribution Networks.” Presented at the 2020 IEEE Power & Energy Society General Meeting (PESGM).

Xu, Z., M. A. Figueiredo, X. Yuan, C. Studer, and T. Goldstein. 2017. “Adaptive Relaxed ADMM: Convergence Theory and Practical Implementation.” Presented at the IEEE Conference on Computer Vision and Pattern Recognition. Honolulu, Hawaii.

Yamashita, D. Y., I. Vechiu, and J.-P. Gaubert. 2020. “A Review of Hierarchical Control for Building Microgrids.” *Renewable and Sustainable Energy Reviews* 118.

Yokota, T., and H. Hontani. 2017. “An Efficient Method for Adapting Step-Size Parameters of Primal-Dual Hybrid Gradient Method in Application to Total Variation Regularization.” Presented at the Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC). Kuala Lumpur, Malaysia.