

Flattening the Duck Curve: A Case for Distributed Decision Making

Rabab Haider*, Giulio Ferro†, Michela Robba†, and Anuradha M. Annaswamy*

*Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, USA

† DIBRIS, University of Genoa, Genoa, Italy

Email: rhaider@mit.edu, giulio.ferro@edu.unige.it, michela.robba@unige.it, aanna@mit.edu

Abstract—The large penetration of renewable resources has resulted in rapidly changing net loads, resulting in the characteristic “duck curve”. The resulting ramping requirements of bulk system resources is an operational challenge. To address this, we propose a distributed optimization framework within which distributed resources located in the distribution grid are coordinated to provide support to the bulk system. We model the power flow of the multi-phase unbalanced distribution grid using a Current Injection (CI) approach, which leverages McCormick Envelope based convex relaxation to render a linear model. We then solve this CI-OPF with an accelerated Proximal Atomic Coordination (PAC) which employs Nesterov type acceleration, termed NST-PAC. We evaluate our distributed approach against a local approach, on a case study of San Francisco, California, using a modified IEEE-34 node network and under a high penetration of solar PV, flexible loads, and battery units. Our distributed approach reduced the ramping requirements of bulk system generators by up to 23%.

Index Terms—Distribution grid, Distributed optimization, Energy storage

I. INTRODUCTION

The push towards decarbonization of the electric grid has seen an explosive growth in renewable energy installation. The state of California generates approximately 25% of its electricity demand from solar resources, resulting in the characteristic “duck curve” in net system load [1]. This new operating condition, where dispatchable bulk resources must quickly meet the large and rapid change in electricity demand, introduces challenges to grid operators. Storage offers one solution, as a flexible and low-inertia resource, however remains too expensive for utility-scale system-wide adoption. Instead, we propose system operators look towards the distribution grid to provide some support. Small-scale consumer owned distributed energy resources (DERs) which include rooftop PV with inverters, demand response, and batteries can provide support to the bulk grid.

Recent works in this area are [2] which proposes a game theoretic framework for dynamic pricing, and [3] which shifts thermal cooling loads of residential units to mitigate duck curve effects. Where these works rely on local approaches with assets locally providing the mitigation mechanisms, our

This material is based upon work partially supported by the MathWorks Mechanical Engineering Fellowship, U.S. Department of Energy under Award Number DE-IA0000025, NSF CPS-DFG Joint award 1932406, and the DIBRIS-UNIGE on institutional funding SEED 2019.

proposed work takes a systems-approach. We leverage distributed optimization to deploy resources over a larger spacial area of the distribution grid to enhance mitigation strategies while suitably accommodating global network constraints.

The main contribution of the paper pertains to the distributed coordination of DERs across a large region of the distribution grid to mitigate the duck curve. This coordination is accomplished by leveraging the variable power factor setting of PV inverters, flexible loads to reduce consumption, and distributed storage devices such as community batteries. This central contribution is realized by formulating the global optimization problem as an optimal power flow, and employing a current-injection (CI) based linear model describing the power physics of the unbalanced distribution grid. The CI-based OPF is solved in a distributed fashion, using a Proximal Atomic Coordination (PAC) approach [4] with Nesterov acceleration, called NST-PAC. We evaluate the performance on a modified IEEE-34 node network, using load and generation data for San Francisco. Our results show the successful mitigation of the duck curve through the coordination of spatially distributed DERs.

In Sections II and III we introduce the CI power flow model for the unbalanced distribution grid, and the accelerated distributed optimization algorithm called NST-PAC. In Section IV we carry out a numerical case study, comparing local and distributed paradigms. We provide conclusions in Section V.

II. CURRENT INJECTION MODEL

The distribution grid is a highly unbalanced network, with many single- and two-phase lines and corresponding loads. As DER penetration increases, many of which are added to single-phase lines, the network will become more unbalanced. In light of this, power flow models which assume balanced flow [5] or those which are valid for only a small range of angle imbalances [6] cannot suitably describe the future grid. Further, the faster timescales of operation due to renewable variability and flexibility requirements need a model which is simple and results in OPF problems which are computationally tractable. For this reason, there is interest in developing scalable linear models for the unbalanced grid, and leveraging distributed computation and algorithms to solve the OPF problem.

In this work, we utilize the Current Injection (CI) model, a linear model for multi-phase unbalanced distribution grids. The CI model takes a similar approach to the Bus Injection

Model, wherein all nodal power injections are modeled as current injections, and all phasor variables are represented in Cartesian coordinates. The 3-phase impedance matrix is used to describe the self and mutual inductance between phases to model the coupling of phases that are common to a distribution grid. The non-convexity of the AC-OPF and the subsequent nonlinearity of SOCP and SDP convexification strategies typically used are avoided in the CI approach by leveraging McCormick Envelope (MCE) based convex relaxation [7] for the bilinear power relations. The MCE uses the convex hull representation of bilinear terms to render a linear OPF model. Suitable pre-processing techniques can be utilized to determine adequate bounds on the nodal voltages and currents to ensure a tight convex relaxation.

A. Problem Formulation

We denote a general distribution network as a graph $\Gamma(\mathcal{N}, \mathcal{E})$, where $\mathcal{N} := \{1, \dots, N\}$ denotes the set of nodes, $\mathcal{E} := \{(m, n)\}$ denotes the set of edges, and each phase is expressed as $\phi \in \mathcal{P}$, $\mathcal{P} = \{a, b, c\}$. The CI-OPF for the network is then written as¹:

$$\min_x f(x) \quad (1a)$$

$$AV = ZI_{\text{flow}} \quad (1b)$$

$$I^R = \text{Re} \left(A^T I_{\text{flow}} \right), \quad I^I = \text{Im} \left(A^T I_{\text{flow}} \right) \quad (1c)$$

$$P_j^\phi = V_j^{\phi, R} I_j^{\phi, R} + V_j^{\phi, I} I_j^{\phi, I} \quad (1d)$$

$$Q_j^\phi = -V_j^{\phi, R} I_j^{\phi, I} + V_j^{\phi, I} I_j^{\phi, R} \quad (1e)$$

$$\underline{P}_j^\phi \leq P_j^\phi \leq \bar{P}_j^\phi, \quad \underline{Q}_j^\phi \leq Q_j^\phi \leq \bar{Q}_j^\phi \quad (1f)$$

$$V_j^{\phi, R} \leq \underline{V}_j^{\phi, R} \leq \bar{V}_j^{\phi, R}, \quad V_j^{\phi, I} \leq \underline{V}_j^{\phi, I} \leq \bar{V}_j^{\phi, I} \quad (1g)$$

$$\underline{I}_j^{\phi, R} \leq I_j^{\phi, R} \leq \bar{I}_j^{\phi, R}, \quad \underline{I}_j^{\phi, I} \leq I_j^{\phi, I} \leq \bar{I}_j^{\phi, I} \quad (1h)$$

where $x = [I^R \ I^I \ V^R \ V^I \ P \ Q \ I_{\text{flow}}^R \ I_{\text{flow}}^I]$ is the decision vector for the CI-OPF problem; I, V, P, Q denote the vector of nodal current injections, voltages, and real/reactive power injections respectively; I_{flow} denotes the vector of line currents; $A \in \mathbb{R}^{3N \times 3N}$ is the 3-phase graph incidence matrix; Z is the system impedance matrix. We use x^R and x^I to denote the real and imaginary components of a complex number x ; overbar \bar{x} and underbar \underline{x} denote the upper and lower limits of a variable x ; $\text{Re}(\cdot)$ and $\text{Im}(\cdot)$ denote the real and imaginary components of a complex number. Constraint (1b) describes the generalized Ohm's law, (1c) describe Kirchhoff's Current Law, and (1d)-(1e) are the definitions of real and reactive power. Constraints (1d)-(1h) are for all nodes $j \in \mathcal{N}$, per each phase $\phi \in \mathcal{P}$, and for all time t . The convex relaxation of the CI-OPF problem through McCormick envelopes adds additional constraints to the problem namely (1h), concerning the minimum and maximum limits of injected current at each phase of each node. The determination of these upper and lower bounds on nodal current is not trivial. An effective heuristic to define these values is presented in [9].

¹For brevity, we omit the McCormick Envelope relaxation of the bilinear terms and the pre-processing step. See [8], [9] for details

B. Model of DERs

We use active sign convention, such that nodal injections are positive. The PV units are modelled as a generator with variable power factor (pf), equipped with a multiphase inverter, where the ratio of P and Q determine the pf setting:

$$P_j^\phi \tan(\cos^{-1}(-pf)) \leq Q_j^\phi \leq P_j^\phi \tan(\cos^{-1}(pf)) \quad (2)$$

Flexible loads are reductions in real power demand (reactive power loads are fixed), and are modelled as a percentage reduction from the forecasted load. The power constraints are modified as $\bar{P}_j(t) = \underline{P}_j(t) * (1 - \alpha_j^{\text{DR}}(t))$, by recalling the use of active sign convention. For inflexible loads, $\underline{P}_j = \bar{P}_j, \bar{P}_j \leq 0$. Prosumers, which are nodes where both load and generation are present, are modelled with additional variables representing both load and generation, P^L and P^G respectively. For prosumers, the inverter is modelled on P^G and Q^G , and load flexibility is modelled on P^L . Variables P^L and P^G are both nonnegative, and the same for reactive power.

$$P_j = P_j^G - P_j^L, \quad P_j^G \geq 0, P_j^L \geq 0 \quad (3a)$$

$$Q_j = Q_j^G - Q_j^L, \quad Q_j^L \geq 0 \quad (3b)$$

$$\bar{P}_j = \bar{P}_j^G - \underline{P}_j^L, \quad \underline{P}_j = \underline{P}_j^G - \bar{P}_j^L \quad (3c)$$

$$\bar{Q}_j = \bar{Q}_j^G - \underline{Q}_j^L, \quad \underline{Q}_j = \underline{Q}_j^G - \bar{Q}_j^L \quad (3d)$$

Battery storage devices are modelled using the power charge and discharge, $P_j^{\text{sc}}(t)$ and $P_j^{\text{sd}}(t)$ respectively for node j and time t . These are nonnegative variables. The state of charge, $b_j(t)$, is calculated as an integral constraint using the actions of the previous period and the initial state of charge, $b_j^0 := b_j(t=0)$. We model charge and discharge efficiencies (η_j^C and η_j^D), self-discharge rate (η_j^{self}), and impose a minimum state of charge (\underline{b}_j) to ensure battery health. We assume all batteries operate at unity power factor, i.e. $Q_j = 0 \quad \forall t$.

$$P_j = \frac{1}{\eta_j^D} P_j^{\text{sd}}(t) - \eta_j^C P_j^{\text{sc}}(t) \quad (4a)$$

$$0 \leq P_j^{\text{sd}} \leq \bar{P}_j^{\text{sd}}, \quad 0 \leq P_j^{\text{sc}} \leq \bar{P}_j^{\text{sc}} \quad (4b)$$

$$b_j(t) = (1 - \eta_j^{\text{self}})b_j(t-1) + \eta_j^C P_j^{\text{sc}}(t) - \frac{1}{\eta_j^D} P_j^{\text{sd}}(t) \quad (4c)$$

$$\underline{b}_j \leq b_j(t) \leq \bar{b}_j \quad (4d)$$

The storage model introduces a dependency on control action in one period to previous periods. The other DERs and loads do not have inter-temporal constraints, so constraints (1)-(3) are simply replicated for each time step t .

III. PAC-BASED DISTRIBUTED IMPLEMENTATION

In this section we introduce a PAC-based distributed optimization algorithm to coordinate a large number of spatially distributed DERs at the grids edge. Consider a global optimization problem composed of equality and inequality constraints, which may be coupled in time

$$\begin{aligned} \min_x & \sum_{i=1}^S f_i(x) \\ \text{s. t. } & Gx = b, \quad Hx \leq d \end{aligned} \quad (5)$$

where $\sum_{i \in S} f_i(x)$ represents the total objective function. Problem (5) can be distributed into $K = \{1, \dots, k, \dots, K\}$ separate coupled optimization problems, denoted as *atoms*. We use a decomposition profile which separates the vector of variables x into two sets: $L = \{L_j, \forall j \in K\}$ and $O = \{O_j, \forall j \in K\}$, which represent the partition of decision variables “owned” and “copied” by atom j . The set of total variables (owned and copied) by an atom is denoted as T . The decomposition profile also separates the constraints into sets owned by each atom, as $C = \{C_j, \forall j \in K\}$. The notion of variables copies are used to satisfy the coupling in constraints and/or objective function. In the context of the CI model, the power physics of the grid (1b)-(1c) result in these coupling constraints. Note that the CI does not have coupling introduced by inequality constraints, however the decomposition can be trivially extended to coupled inequality constraints. Using the decomposition profile, we obtain:

$$\begin{aligned} & \min_{a_j} \sum_{j \in K} f_j(a_j) \\ \text{subj. to: } & \begin{cases} G_j a_j = b_j, & \text{for all } j \in K \\ H_j a_j \leq d_j, & \text{for all } j \in K \\ B_j a = 0, & \text{for all } j \in K \end{cases} \end{aligned} \quad (6)$$

where a_j is atom j ’s variables (both owned and copied), $f_j(a_j)$ is the atomic objective function, and G_j , b_j , H_j , and d_j represent the submatrix or subvector of G , b , H , and d respectively. Finally, B is in incidence matrix over the owned and copied atomic variables, defined as

$$B_i^m \triangleq \begin{cases} -1, & \text{if } i \text{ is ‘owned’ and } m \text{ a related ‘copy’} \\ 1, & \text{if } m \text{ is ‘owned’ and } i \text{ a related ‘copy’} \\ 0, & \text{otherwise} \end{cases}$$

We then use B_j (B^j) to denote the relevant incoming (outgoing) edges of the directed graph for atom- j . To fully parallelize the optimization, we introduce *coordination constraints*, which must be satisfied for every atom. These require all atomic copied variables in a given j th atom to equal the value of their corresponding owned in i th atom, $i \neq j$:

$$B_j a = 0 \quad \forall j \in K \quad (7)$$

A. An Accelerated Algorithm: NST-PAC

We next present an accelerated variant of the PAC algorithm in [4], which includes time-varying gains and Nesterov type acceleration [10], called NST-PAC. The NST-PAC is a primal-dual method with ℓ_2 and proximal regularization, Nesterov type acceleration for both primal and dual variables, and privacy-preserving features. We begin by forming the atomic Lagrangian function:

$$\begin{aligned} \mathcal{L}(a, \mu, \nu) &= \sum_{j \in K} \left[f_j(a_j) + \mu_j^T (G_j a_j - b_j) + \nu^T B_j a \right] \\ &= \sum_{j \in K} \left[f_j(a_j) + \mu_j^T (G_j a_j - b_j) + \nu^T B^j a_j \right] \\ &\triangleq \sum_{j \in K} \mathcal{L}_j(a_j, \mu_j, \nu). \end{aligned} \quad (8)$$

The algorithm is carried out as below:

$$\begin{aligned} a_j[\tau + 1] &= \underset{a_j \in \mathbb{R}^{|T_j|}}{\operatorname{argmin}} \left\{ \mathcal{L}_j(a_j, \hat{\mu}_j[\tau], \hat{\nu}[\tau]) + \frac{\rho_j \gamma_j}{2} \|G_j a_j - b_j\|_2^2 \right. \\ &\quad \left. + \frac{\rho_j \gamma_j}{2} \|B_j a_j\|_2^2 + \frac{1}{2\rho_j} \|a_j - a_j[\tau]\|_2^2 \right\}, \end{aligned} \quad (9)$$

$$a_j[\tau + 1] = a_j[\tau + 1] + \alpha_j[\tau + 1](a_j[\tau + 1] - a_j[\tau]) \quad (10)$$

$$\mu_j[\tau + 1] = \hat{\mu}_j[\tau] + \rho_j \gamma_j (G_j \hat{a}_j[\tau + 1] - b_j) \quad (11)$$

$$\hat{\mu}_j[\tau + 1] = \mu_j[\tau + 1] + \phi_j[\tau + 1](\mu_j[\tau + 1] - \mu_j[\tau]) \quad (12)$$

Communicate \hat{a}_j for all $j \in [K]$ with neighbors (13)

$$\nu_j[\tau + 1] = \hat{\nu}_j[\tau] + \rho_j \gamma_j B_j \hat{a}_j[\tau + 1] \quad (14)$$

$$\hat{\nu}_j[\tau + 1] = \nu_j[\tau + 1] + \theta_j[\tau + 1](\nu_j[\tau + 1] - \nu_j[\tau]) \quad (15)$$

Communicate $\hat{\nu}_j$ for all $j \in [K]$ with neighbors (16)

where ρ_j , γ_j are atom-varying over-relaxation and step-size parameters, respectively. The proposed NST-PAC uses ℓ_2 regularization terms rather than the prox-linear variant in PAC. Further, both primal and dual variables are accelerated using Nesterov type acceleration, to speed up convergence. Further, we extend the privacy-preserving feature of the PAC algorithm to both the primal and dual variables, by using three iteration-varying and atom specific parameters for the accelerated terms, $\alpha_j[\tau] > \alpha_j^{\min}$, $\phi_j[\tau] > \phi_j^{\min}$ and $\theta_j[\tau] > \theta_j^{\min}$. In the original algorithm privacy is kept only for dual variables. A detailed analysis of convergence rate, communication and computational complexity, and privacy are provided in [4].

IV. CASE STUDY

We consider a case study of San Francisco, California, using the IEEE-34 node network as a proxy for the distribution grid. The load data from the IEEE datasheet serves as the daily average load, and the 24-hr load profiles are obtained from the ODEI dataset from NREL for the Typical Meteorological Year [11]. All loads are assumed to have a constant power factor of 0.95. The network loads are classified as residential or commercial loads based on the size of the load, by matching the load levels of the IEEE-34 network with the TMY data. Commercial loads include retail space, small and medium office buildings, primary school, medium and large restaurants, and a hospital. The network is modified to include DERs which include clusters of rooftop photovoltaic (PV) units, flexible loads, and three battery storage units. The penetration of PV is 38%, as measured by the ratio of nameplate capacity to average system load. This high DER penetration scenario is a reasonable projection given the RPS initiatives in California. Each PV unit is assumed to be equipped with an inverter with corresponding power electronic control, which can be operated at variable power factor in the range of 0.8 to 1. PV curtailment is not considered. The flexible loads are modelled as typical residential cooling loads for California [12]. Variations in nodal demand response are obtained by shifting the baseline profile obtained from [12] in time and space, with both following zero-mean Gaussian distributions with variances of 0.075 and 0.1 respectively. The three battery

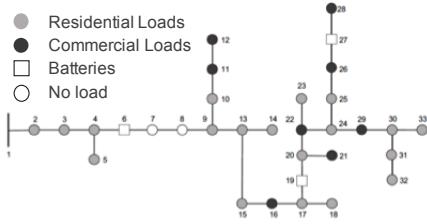


Fig. 1: Topology of IEEE-34 network. PV units and flexible loads are present throughout the network.

units are a 450kW-120kWh community unit² at node 6, a cluster of 40 Tesla Powerwall+ batteries (each 13.5kW-5kWh) at node 19, and a 800kW-185kWh hospital unit at node 27. Any load not met by local generation (or storage) is assumed to be served by the bulk grid at the point of common coupling (PCC), at node 1 in the network ($j = 1$). Figure 1 shows the network topology.

A. Scenarios

We consider three scenarios, as below:

- **Scenario A: Baseline.** All PV inverters operate at unity pf, and batteries and flexible loads are not present.
- **Scenario B: Local control.** Each DER owner operates its devices and manages its loads.
- **Scenario C: Distributed control.** All devices and loads are coordinated using the NST-PAC algorithm.

Scenario A quite trivially is the characteristic duck curve, where the high PV generation at unity power factor results in large ramping requirements of transmission-level generators. Scenario B is a local approach, where each agent will minimize its peak load throughout the day. This serves as a very rough approximation of reducing the ramping requirements of bulk resources, by noticing that the largest ramp typically coincides with the peak demand in the evening. This action can be motivated by the fact that consumers can be charged based on their peak energy consumption. To leverage the capabilities of the storage devices, neighbouring nodes are clustered with the battery. Residential loads (and corresponding DERs) at nodes 3, 4, and 5, share the community battery at node 6. Residential loads at node 20 and the primary school at node 21 share the cluster of Tesla Powerwall+ batteries at node 19. The hospital at node 26 is assumed to own and operate the battery at node 27. All remaining nodes are treated as independent agents. After clustering, there are a total of 26 agents in the network, each managing its own consumption and generation, to minimize its peak load. The multi-period optimization problem solved by each agent is a simple power balance, where any load in excess of local generation is assumed to be served by the bulk system. The power physics between nodes within a cluster is not modelled. The objective function is $f_{\text{local}}(y) = \max_t \{-P(t)\}$, where $y = [P, Q, P^L, P^G, Q^L, Q^G, P^{\text{sd}}, P^{\text{sc}}]$.

²Modeled on the Ellenbrook unit from the PowerBank trial, Australia [13]

Scenario C is the PAC-based approach which requires coordination between the devices. This approach accommodates system-level constraints including grid power physics, and minimizes the ramping requirements at the PCC. In Scenario C, the PAC agents solve the CI-based multi-phase unbalanced OPF, to minimize the objective function $f(x)$ which minimizes the difference in power supplied by the PCC from one hour to the next. The function is $f_{\text{PAC}}(x) = \left\| \sum_{\phi \in \mathcal{P}} P_1^\phi(t) - P_1^\phi(t-1) \right\|$.

B. Results and Discussion

Each scenario was simulated on the IEEE-34 node network. Figure 3 shows the net load served by the bulk system. The load curve for Scenario A shows the characteristic high ramps down and up when solar generation begins and ends. Scenario B provides minimal improvement, reducing the load for most hours of the day, while Scenario C effectively leverages the DERs to reduce the ramping requirements. Note that the objective function reduces the hour-to-hour change in load, and to do so, increases load during the hours of peak PV generation (roughly 10am to 4pm) to charge the batteries, which are then discharged in the late evening to reduce the net load. The magnitude of hour-to-hour ramping is shown in Figure 4. Notably, the local optimization of Scenario B, which minimizes the peak load throughout the day as a proxy for minimizing the load ramp in the evening, is not able to suitably reduce the ramping requirement. The distributed optimization approach, on the other hand, is able to leverage system-wide information and coordinate the DERs to provide grid-level support, as needed by the bulk system. This coordinated approach is able to reduce ramping requirements throughout the day, with a 23% reduction in ramping requirements at the 4pm peak. Table I presents the total ramping reduction for Scenarios B and C, as compared to the baseline in A, and the computational run times. As expected, the proposed distributed coordination significantly outperforms the local approach, with 28% reduction in ramping required. The local approach is unable to provide any reduction. Unsurprisingly, the local approach takes less time to reach a decision (albeit an inferior one), while the distributed approach takes considerably longer (completing 1000 iterations). However, computational time of 20s is still well within the acceptable time-frame for decision making, and enables DERs to provide bulk-level support.

We next investigate the impact of the battery units' initial state of charge. We run each Scenario for (1) Minimum SOC where initial capacity is at 45, 0, and 160 kWh; (2) Mid SOC where initial capacity is at 120, 400, and 400 kWh; and (3) Full SOC where initial capacity is at 450, 540, and 800 kWh, respectively for batteries at nodes 6, 19, and 27. Figure 2 plots the SOC of each battery unit for the three cases. The usage pattern in each of the cases is very similar, with charging in the early morning and afternoon to build up storage capacity for the evening. Interestingly, the final SOC are non-zero for the three cases, and are quite similar, suggesting the batteries may be able to retain a higher minimum charge to be used as backup power in the case of emergencies. The batteries

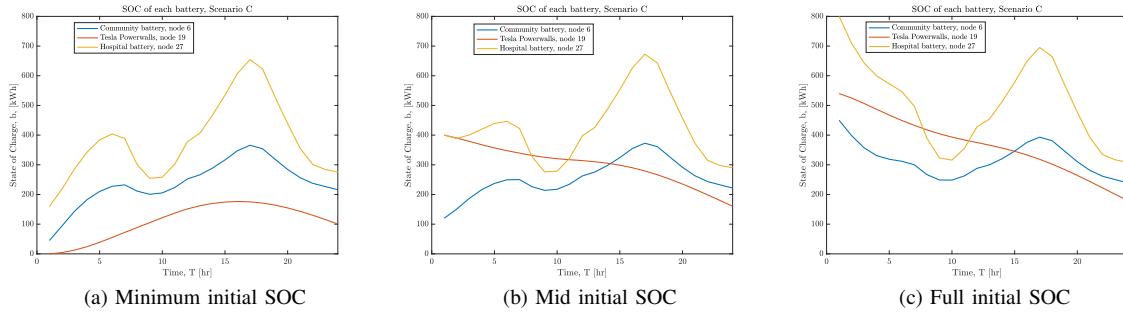


Fig. 2: State of charge of each storage device for Scenario C, with different initial state of charge for each battery.

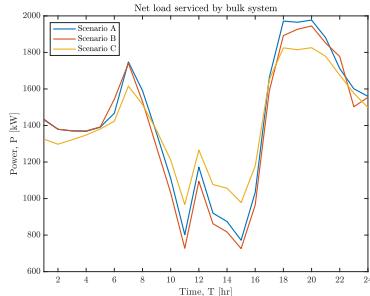


Fig. 3: Net load serviced by the bulk system, for each scenario. These plots do not include power loss over lines.

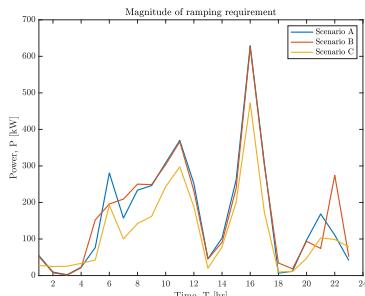


Fig. 4: Magnitude of ramping required by bulk system generators for each scenario. provide flexibility to increase or decrease load throughout the day, as required by the grid. The plots for net load are quite similar for all three cases, and so have been omitted.

V. CONCLUSION

We presented and evaluated a distributed approach to coordinate DERs to provide services to the bulk system. We leverage a CI-based linear model of the unbalanced grid, and an accelerated PAC-based algorithm called NST-PAC. Our case study on a modified IEEE-34 node network shows how distributed techniques can leverage information from different

resources to successfully mitigate the duck curve, reducing ramping requirements of bulk system generators by up to 23%. This framework can be extended to include electric vehicles, by modeling the vehicle-to-grid capabilities and corresponding cost of battery cycling and lifetime degradation in the optimization problem. Using such an approach throughout the distribution grid can reduce the challenges of the new operating condition resulting from high solar and renewable penetration, reduce system costs, and improve renewable integration. Future work will concern the development of faster distributed algorithms applicable to nonconvex costs and constraints, with discrete and continuous controls.

REFERENCES

- [1] CAISO, “California ISO demand response and energy efficiency roadmap: Maximizing preferred resources,” 2013.
- [2] M. Sheha, K. Mohammadi, and K. Powell, “Solving the duck curve in a smart grid environment using a non-cooperative game theory and dynamic pricing profiles,” *Energy Conversion and Management*, vol. 220, p. 113102, 2020.
- [3] I. Calero, C. A. Cañizares, K. Bhattacharya, and R. Baldick, “Duck-curve mitigation in power grids with high penetration of pv generation,” *IEEE Transactions on Smart Grid*, vol. 13, no. 1, pp. 314–329, 2021.
- [4] J. Romvary, G. Ferro, R. Haider, and A. M. Annaswamy, “A distributed proximal atomic coordination algorithm,” *IEEE Transactions on Automatic Control*, 2021. doi: 10.1109/TAC.2021.3053907.
- [5] L. Gan, N. Li, U. Topcu, and S. H. Low, “Exact convex relaxation of optimal power flow in radial networks,” *IEEE Transactions on Automatic Control*, vol. 60, no. 1, pp. 72–87, 2014.
- [6] B. A. Robbins and A. D. Domínguez-García, “Optimal reactive power dispatch for voltage regulation in unbalanced distribution systems,” *IEEE Transactions on Power Systems*, vol. 31, no. 4, pp. 2903–2913, 2015.
- [7] G. P. McCormick, “Computability of global solutions to factorable nonconvex programs: Part i — convex underestimating problems,” *Mathematical Programming*, vol. 10, no. 1, p. 147–175, 1976.
- [8] G. Ferro, M. Robba, D. D’Achiardi, R. Haider, and A. M. Annaswamy, “A distributed approach to the optimal power flow problem for unbalanced and mesh networks,” *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 13287–13292, 2020. 21th IFAC World Congress.
- [9] G. Ferro, *Competitive and Cooperative Approaches to the Balancing Market in Distribution Grids*. PhD thesis, Università degli studi di Genova, 2020.
- [10] Y. E. Nesterov, “A method for solving the convex programming problem with convergence rate $O(1/k^2)$,” in *Dokl. akad. nauk Sssr*, vol. 269, pp. 543–547, 1983.
- [11] National Renewable Energy Laboratory, “Commercial and residential hourly load profiles for all TMY3 locations in the United States [data set],” 2014. doi:10.25984/1788456.
- [12] D. Olsen, M. Sohn, M. A. Piette, and S. Kilicotte, “Demand response availability profiles for California in the year 2020,” *LBNL Report*, 2014.
- [13] P. M. Tagabe, “The power of community batteries: bringing renewables to the masses,” *Energy Magazine*, Aug 2020.

	A (Baseline)	B (Local)	C (Proposed)
Total ramping need (kW)	3923.7	3941.0	2839.4
Ramping reduction	-	-0.44%	27.63%
Mean run time per agent	-	0.0947s	16.97s

TABLE I: Summary of results for Scenarios A thru C. Scenario A is baseline with no decision making, so no ramping reduction or computational time to report.