Methodology for comparing the performance of DER coordination schemes in providing frequency regulation

Hani Mavalizadeh, *Student Member, IEEE*, Mads R. Almassalkhi, *Senior Member, IEEE*

***Abstract*—In this paper, we illustrate a novel methodology for comparing and quantifying the performance of different distributed energy resources (DER) schemes' ability to deliver frequency regulation services across a number of salient criteria. The schemes considered include (bottom-up) packetized energy management, fitness-based methods, and (direct load control) optimization-based methods. The criteria of interest include tracking performance, scalability of communication, scalability of computation, device availability, ability to maintain consumer quality of service (CQoS) relative to delivered hot water temperatures, and impact on the device quality of service (DQoS) such as average cycling rates. Moreover, we augment the fitness-based method with an ability to estimate the fitness values dynamically which significantly reduces the communication burden while maintaining the tracking capability. Finally, the simulations and corresponding comparisons are based on a representative subset of PJM's historical Reg-D data.**

***Index Terms*—Distributed energy resources, fitness-based coordination, packetized energy management, flexibility, frequency regulation**

I. Introduction

Public opinion and policies concerning efforts to mitigate climate change are driving increased penetration of renewable generation. Integrating renewable energy sources into the electric grid while maintaining reliability is a fundamental power engineering challenge that will require large-scale deployment of grid-side and demand-side flexibility. New advances in sensor technology together with low-cost edge computing and connectivity enable DER coordinators to regulate DERs remotely to respond to the needs of the grid and leverage market opportunities while satisfying customer usage [1]. This makes demand-side management a viable option for ancillary services, such as frequency regulation [2] or fast frequency response [3]. As a result, much research is focused to study the potential uses of DER coordination in modern power system operations in recent years.

For the DER coordination schemes to be valuable for system operations, they should be capable of coordinating thousands of kW-scale, flexible electric loads such as electric vehicles (EVs), batteries, and thermostatically controlled loads (TCLs) [4]. At this scale, the roles of computation, communication, and data management requirements are critical. In addition, most of the DER coordination schemes require that some data is shared between consumers and a coordinator. This raises data privacy [5], as well as, cyber-security concerns [6]. In fact, if consumers are not convinced that their information is safe and secure, the participation rate will drop, rendering the whole scheme unviable [7], [8]. Thus, DER coordination architecture requires careful analysis and design.

Both authors are with the Department of Electrical and Biomedical Engineering, University of Vermont, Burlington, VT 05405 and graciously acknowledge support from the U.S. National Science Foundation (NSF) Award ECCS-2047306 (e-mails: {hmavalz,malmassa}@uvm.edu).

DER coordination schemes can generally be categorized into direct/top-down and indirect/bottom-up architectures. In direct control schemes, a central coordinator has full access to and control over all DER information and actuation. While direct control potentially leads to good performance, the large computation and communication burden raises concerns about scalability. In indirect control schemes, the decisions are made at each DER based on local measurements, which reduces communication and computation overhead, but raises concerns about tracking performance

Another key factor in DER coordination architectures is the underlying data availability and communication requirements [9]. Indirect schemes can be further classified based on the level of information shared between consumer and coordinator: i) mediated coordination: coordinator collects DER information, ii) bilateral coordination: DERs communicate with each other, and iii) implicit coordination: individual DER information is not shared [9]. While higher levels of information sharing improve the performance of the coordination scheme, it can lead to privacy and security issues.

In this paper, key metrics that underpin DER coordination architectures are presented and illustrated for three different DER control architectures: i) a bottom-up device-driven scheme called packetized energy management (PEM) [10]; ii) a fitness-based prioritization scheme [11], and iii) an optimization-based direct scheduling scheme [12]. The main contributions of this paper are listed below:

- Quantitative methodology is proposed for holistically analyzing the performance of DER coordination schemes across a set of proposed salient and practical metrics.
- A fitness-based DER coordination scheme is specifically extended by enabling the DER coordinator to dynamically update the DER fleet's fitness values, which permits significantly lower communication burden and DQoS without negatively impacting the ability to deliver grid services and CQoS.
- The real-time, cyber-enabled DER simulation platform from [13] is extended to incorporate two more DER coordination schemes. This improved platform is utilized in the simulation of a fleet of 1000 electric water heaters (EWHs) to illustrate the quantitative methodology for different DER coordination schemes from the literature.

The remainder of the paper is organized as follows: In Section II, different coordination schemes are briefly explained. The methodology for comparing different schemes is presented in Section III The simulation-based analysis is presented in Section IV while Section V concludes the paper.

II. DER COORDINATION SCHEMES

Fleets of thermostatically controlled loads (TCLs) are used by the DER coordinator to provide frequency regulation. Thus, we first present a model of TCL. Second, we present three different coordination schemes from the literature.

A. Thermostatically controlled load model

Consider a first-order, difference equation modeling the evolution of the tank temperature, z_n, of an electric water heater (EWH):

$$z_n[k+1] = z_n[k] + \Delta t\left(\frac{\eta_n P_n^{\text{rate}}}{c\rho L_n} s_n[k] - \frac{z_n[k]-T_{\text{amb}}}{\tau_n} - \frac{z_n[k]-T_{\text{inlet}}}{60L}\omega_n[k]\right), \tag{1}$$

where c, ρ, and L are the specific heat capacity, water density at 50 degrees, and the tank capacity, respectively, while P_n^{rate} is the power consumed by the heating element and η_n the efficiency of delivering power to the water. The binary ON(1)/OFF(0) variable is denoted $s_n[k]$ while insulation losses are defined by τ_n. Finally, T_{amb},T_{inlet} are the ambient and inlet temperatures, respectively, while $\omega_n[k]$ is the stochastic water withdrawal rate and Δt is the sampling time.

The goal of any DER coordinator managing N TCLs is then to maintain CQoS (i.e., keep $z_n[k]$ close to a customer's desired set-point $z_n^{\text{set}} \in (z_n^{\min}, z_n^{\max}) \forall n$) while tracking reference signal in aggregate: minimize $||P_{\text{ref}}[k] - \sum_{n=1}^{N} P_n^{\text{rate}} s_n[k]||_2$. In doing so, a coordinator will need to cycle devices on/off while considering possible device lock-on/off constraints on device operations (impacting DQoS) and rely on computing and communications to be responsive to changes in the reference signal. The next three subsections will examine three different DER coordination schemes that attempt to achieve these goals.

B. Optimization-based method

A naive approach to achieve the coordinator's objectives is to, at each time step, schedule the on/off state of TCLs using optimization-based methods:

$$\min_{s_n[k]} \quad w_1\epsilon + w_2 \sum_{n=1}^{N} (z_n[k+1] - z_n^{\text{set}})^2 \tag{2a}$$

$$\text{s. t.} \quad z_n^{\min} \le z_n[k] \le z_n^{\max} \quad \forall n \tag{2b}$$

$$|P_{\text{ref}}[k] - \sum_{n=1}^{N} s_n[k] P_n^{\text{rate}}| \le \epsilon \tag{2c}$$

$$s_n[k] \in \{0,1\} \tag{2d}$$

ϵ is chosen smaller than P_n^{rate}. The method assumes full knowledge of the fleet's dynamic state. This means that for every time step, all of the devices must send their temperature data to the coordinator. This results in a large communication overhead, which challenges implementation at scale. In addition, the cycling of devices is not captured in the formulation, which should result in low DQoS. In this paper, we consider two versions of (2):

(i) `Opti`(0): Allow DERs to cycle each time-step.
(ii) `Opti`(a): Limit cycling by locking the DER state $s_n[k]$ for a minutes every time we transition.

Since the optimization-based method does not look ahead more than one time step, we expect that locking will serve to simplify the problem (fewer decision variables), reduce cycling (and improve DQoS), and reduce communication overhead (since we do not need state info for locked TCLs), but at the cost of worse tracking and more CQoS violations.

C. Packetized energy management

Packetized energy management (PEM) is a bottom-up/indirect DER coordination scheme that coordinates flexible loads in real-time to harness their potential for different grid services [14]. The scheme uses the concept of energy packets, which is energy delivered as a fixed amount of power for a fixed duration (e.g., 5kW for 3 minutes). In PEM, DERs request energy packets using a probabilistic request rule based on dynamic state $z_n[k]$:

$$Pr(z_n[k]) = 1 - e^{-\mu(z_n[k])\Delta t}, \tag{3}$$

where μ is a function that maps $z_n[k]$ to a mean-time-to-request (or MTTR) as follows:

$$\mu(z_n[k]) = \begin{cases} 0, & z_n[k] \ge z_n^{\max} \\ m_R \frac{z_n^{\max} - z_n[k]}{z_n[k] - z_n^{\min}} \cdot \frac{z_n^{\text{set}} - z_n^{\min}}{z_n^{\max} - z_n^{\text{set}}}, & z_n^{\min} < z_n[k] < z_n^{\max} \\ \infty, & z_n^{\min} \ge z_n[k] \end{cases} \tag{4}$$

with design parameter, $m_R > 0$ chosen to define the inverse of MTTR for $z_n[k] = z_n^{\text{set}}$ The design is aimed at driving the device's dynamic state close to the set point: the lower the $z_n[k]$, the lower the MTTR, which increases the probability of a request being sent to the coordinator, which increases the likelihood of getting access to an energy packet, which drives $z_n[k]$ back up. The requests are sent to a coordinator asynchronously, and the coordinator accepts/rejects requests based on the market or market reference signal. In addition, to guarantee CQoS, PEM is equipped with an opt-out mechanism which permits the DERs to temporarily exit PEM when $z_n[k] \notin [z_n^{\min}, z_n^{\max}]$ and return to conventional operation (e.g., heating water) until $z_n[k] \in (z_n^{\min}, z_n^{\max})$.

In this comparison, two versions of PEM are considered:

(i) `PEM`(a, m_R): fixed packet length a and MTTR parameter m_R for all DERs in fleet.
(ii) `PEM`$([\underline{a}, \bar{a}], m_R)$: every accepted request has random (and uniformly distributed) packet duration, $a \sim U[\underline{a}, \bar{a}]$, [2].

Next, we present and extend a different indirect method called the fitness-based method, which, unlike PEM, keeps a list of device priorities (i.e., ranked by their fitness to turn on and off) and then selects the fittest devices to actuate to minimize tracking error.

D. Fitness-based method

In the fitness-based method, each DER calculates fitness values based on its local dynamic state, $z_n[k]$, its operating state, $s_n[k]$, and its availability of response (e.g., is device locked, opted out, or otherwise unavailable) [15]. The coordinator then receives fitness values from DERs and sorts the devices based on their fitness values. This simplifies the coordinator's process of actively determining which devices to turn on/off as the power reference signal changes.

Thus, this method depends on device computing and the coordinator ranking a set of fitness values. There are different ways to compute fitness values based on available (local) device information: $z_n[k], s_n[k]$, including time since the last on/off transition.

1) DER fitness values: The fitness values for device n can be defined based only on the dynamic state:

$$F_n^{\text{ON}}[k] = 1 - \exp\left(-\frac{z_n^{\max} - z_n[k]}{z_n[k] - z_n^{\min}}\right) \in [0,1] \quad (5a)$$

$$F_n^{\text{OFF}}[k] = 1 - F_n^{\text{ON}} \in [0,1]. \quad (5b)$$

For example, for EWHs, if a device's temperature is low, then it has a high fitness value for turning ON (F_n^{ON}) and a low fitness value for turning OFF (F_n^{OFF}). This prioritizes EWHs with lower temperatures (high fitness) to be selected for turning on if the reference signal increases. Similar to PEM, an internal opt-out mechanism is incorporated in the scheme to ensure that the temperature comfort range ($z_n^{\min}$,$z_n^{\max}$) is not violated. In addition, when a device cycles, it may be locked on/off for a certain duration before it can cycle again. This locked device behavior reduces cycling (i.e., improves DQoS) but also limits the availability of DERs.

Thus, to capture the availability and prioritize DQoS, a device can augment its fitness value with a DQoS term in the fitness function, as shown below:

$$F_n^{\text{ON}}[k] = \left(1 - \exp\left(-\frac{z_n^{\max} - z_n[k]}{z_n[k] - z_n^{\min}}\right)\right)\left(1 - \exp(-\tau_n^{\text{off}}[k])\right) \quad (6a)$$

$$F_n^{\text{OFF}}[k] = \exp\left(-\frac{z_n^{\max}[k] - z_n[k]}{z_n[k] - z_n^{\min}}\right)\left(1 - \exp(-\tau_n^{\text{on}}[k])\right) \quad (6b)$$

where $\tau_n^{\text{off/on}}$ is the time elapsed (in hours) since the last off/on transition. The DER will then send the coordinator its fitness values based on either (5) or (6), which is used in actively coordinating the DER fleet and described next.

2) Fitness-based coordination and estimation: To achieve the desired coordinator objectives (tracking and CQoS), the coordinator forms separate queues for on and off fitness values. When aggregate DER fleet power, $\sum_{n=1}^{N} s_n[k]P_n^{\text{rate}}$, is larger than the reference signal, $P_{\text{ref}}[k]$, the coordinator will select the devices with the highest $F_n^{\text{OFF}}[k]$ to turn off until $|P_{\text{ref}}[k] - \sum_{n=1}^{N} s_n[k]P_n^{\text{rate}}| \leq \epsilon$. The case when aggregate fleet power is lower than the reference signal is similar with the coordinator selecting the highest $F_n^{\text{ON}}[k]$ devices to turn on. A key challenge with the fitness-based method is the communication overhead associated with the coordinator keeping a fleet's fitness values up-to-date. That is, since TCLs compute their fitness values based on their dynamic state, which is a function of background demand, a device's fitness value may change significantly over a period of 5-15 minutes. This means that devices perceived by the coordinator to have a high/medium fitness value, may in-fact no longer be "fit" for coordination and cause unexpected tracking errors or CQoS challenges. To address this challenge, we augment the fitness-based coordinator with a simple dynamic estimate of fitness values based on historical fitness data available to the coordinator. Thus, the dynamic fitness estimation should provide the coordinator with a more accurate estimate of the state of the fitness queues and improve CQoS. In improving CQoS, we expect lower opt-outs and more accurate tracking at an equivalent or lower communication overhead.

Thus, the coordinator is able to estimate the evolution of each device's ON and OFF fitness values with the following simple model:

$$F_{\text{est},n}^{\text{ON}}[k+1] = F_{\text{est},n}^{\text{ON}}[k] + \alpha_n^{\text{ON}}, \quad F_{\text{est},n}^{\text{ON}}[0] = F_n^{\text{ON}}[0] \quad (7)$$

where constant parameter α_n^{ON} is obtained from the coordinator's historical data on device n's fitness values using linear regression. The case of $F_{\text{est},n}^{\text{OFF}}$ and α_n^{ON} is identical. In the interval between devices updating their fitness value, the coordinator uses the dynamic estimate of fitness. Every 5-15 minutes, DERs update their fitness value based on the measured z_n using equation (5) which resets the fitness values at coordinator to the actual ones.

In this paper, we consider a different version of the fitness-based method, `Fit`(a,b,C), where a, b, and C represents the devices' update rate for fitness values (mins), devices' locked out duration after cycling (mins), and whether coordinator employs dynamic fitness estimation or not (i.e., C=E means *with* estimation).

Finally, Fig.1 summarizes the information flow for the aforementioned schemes, highlighting that they all represent different feedback control schemes that are each implemented in the real-time DER simulation platform from [13]. The green, red, and the blue text describes the types of information shared in `Opti`, `Fit`, and `PEM` schemes, respectively. In the next section, we briefly outline the methodology for quantitatively comparing DER coordination schemes.

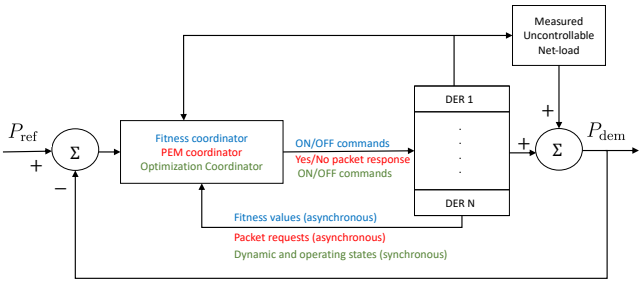

Fig. 1: Feedback control system for different coordination schemes.

III. Methodology for Comparing DER Schemes

To compare DER schemes, we have designed a set of metrics that are relevant for practical implementation and technical evaluation. To quantify the performance of any DER schemes in providing frequency regulation, consider N_s 1-hour samples of PJM's historical Reg-D data set, where each sample $i \in \{1,...,N_s\}$ represents $w_i \in [0,1]$ proportion of the Reg-D data-set and $\sum_i w_i = 1$. Specifically, the metrics used to evaluate the performance of each scheme are denoted as follows for each 1-hour Reg-D sample i:

1) **Device QoS (DQoS):** $M_{4,i} = 1 - X_{4,i}/\max_i\{X_{4,i}\}$, where $X_{4,i}$ is the total number of DER cycles across fleet.
2) **Consumer QoS (CQoS):** $M_{3,i} = 1 - X_{3,i}/\max_i\{X_{3,i}\}$, where $X_{3,i}$ is the number of devices that experience opt-out.
3) **Scalability of communications (SoComm):** $M_{5,i} = 1 - X_{5,i}/\max_i\{X_{5,i}\}$, where $X_{5,i}$ is the average communication in kilobits per second (kbps) per device sent to/from the coordinator from/to DERs in one hour.
4) **Tracking accuracy:** $M_{1,i} = 1 - X_{1,i}/\max_i\{X_{1,i}\}$, where $X_{1,i} = \sqrt{\frac{1}{K}\sum_{k=1}^{K}\left(P_{\text{ref},i}[k] - \sum_{n=1}^{N} P_n^{\text{rate}} s_n[k]\right)^2}$, where $K = 1800$ is the number of time-steps in Reg-D sample i.

5) **PJM composite score:** $M_2 \in [0,1]$ is PJM's formula and includes accuracy, precision, and delay component defined in [2].
6) **Scalability of computing (SoComp):** $M_{6,i} = 1 - X_{6,i}/\max_i\{X_{6,i}\}$, where $X_{6,i}$ is the total processing time used by coordinator.
7) **Device availability:** $M_{7,i} \in [0,1]$ is defined as the average fraction of available devices [16], i.e., devices that are not locked nor opted out.

The above metrics are used to capture different facets of DER coordination. Because of trade-offs in control and communications, no single method will dominate across all metrics. Thus, a utility or aggregator must weigh the metrics based on their preferences. Now, for each 1-hour Reg-D sample $i \in \{1,...,N_s\}$, consider metric j, M_{ji}, for $j = 1,\dots,N_M$. Then, a DER coordination scheme's weighted average metric j is $\overline{M}_j = \sum_{i=1}^{N_s} M_{ji} w_i \ \forall j = 1,...,N_M$. In addition, to further characterize how each scheme performs across the set of N_s samples, we also consider the worst-case performance, $\underline{M}_j = \min_i\{M_{ij}\}$ for all N_M metrics. In the next section, we illustrate the methodology by quantifying the weighted mean and worst-case performance of three different classes of DER coordination across a representative subset of PJM's 8760-hour Reg-D data-set.

IV. Simulation results

In this section, a simulation-based illustration of the methodology for comparing DER schemes is presented. The simulations consider sixteen representative 1-hour Reg-D samples of AGC. The process for selecting the representative samples (and their weights w_i) is briefly described next.

A. Determining representative 1-hour AGC samples

In this section, we select a representative subset of the PJM Reg-D annual (8760 samples) data set by using K-mean clustering. This method partitions the PJM data set into K sub-sets from which we can select representative "centroids" and appropriate weights (related to relative sizes of the sub-sets).

In this work, for each 1-hour Reg-D sample, the sample's average and pegging amount (which is the number of instances in the sample where the reference signal is at ± 1) are used in a K-mean clustering algorithm [17] to find the representative subset of the sample. Based on applying the K-mean algorithm and sweeping across a range of K, it was found that $N_s = 16$ clusters were optimal. Interestingly, the resulting 16 samples represent eight months of the year, all days of the week, and nine day hours. In the next section, simulation-based analysis is used to illustrate the methodology for comparing DER coordination schemes across these N_s samples.

B. Simulation case study

The results are shown next and leverage and extend the real-time, cyber-enabled DER simulation platform, which is detailed in [13]. Specifically, this paper augments the platform with the optimization and fitness-based methods to enable comparisons of DER coordination methods. The simulation platform uses Python 3.9 for the simulation of diverse DER fleets. Local DER control logic and model is implemented in C/C++ based on simplified first-order state of charge dynamics. Gurobi 9.0.1 solver is used for the optimization-based method. The DER coordinator is implemented in Python using requests module and an open-source event-driven networking engine called Twisted. Specifically, Hypertext Transfer Protocol (HTTP) which is a standard protocol for asynchronous communication between Web servers and clients, is used to exchange messages between DERs and coordinator. The platform uses communication mechanisms that closely align with those used by IoT-enabled devices in real-world demand dispatch. Specifically, the simulation platform enables a real-time simulation of 1000 EWHs, each of which has a 4.5 kW power rating, and the DER coordinator responds, in aggregate, to a Reg-D power reference signal that is updated every 2 seconds with a baseline power value equal to 400 kW and an amplitude (capacity) of $\pm$200 kW. The results for mean and worst performance are summarized in Table I where *avg* and *worst* refer to the average and worst values of $X_{j,i}$ across $i = 1,...,N_s$. To better illustrate the results, in Fig. 2, $\overline{M}_j$ is compared for seven DER coordination schemes. In addition, in Fig. 3, the worst performance across all representative 1-hour samples is compared.

By comparing $\texttt{Fit}(15,3,\text{NE})$ and $\texttt{Fit}(15,3,\text{E})$ in table I it can be seen that CQoS is improved significantly when estimation is added to fitness based method (i.e., (7)). On the other hand, by comparing $\texttt{Fit}(3,3,\text{NE})$ and $\texttt{Fit}(3,3,\text{E})$, very small difference is seen. This shows that the estimation method is more useful when the fitness update time. When the update time is small, the error in fitness values is insignificant leading to a small impact on estimation. As expected, by increasing the update time, CQoS decreases due to an increase in the number of opt-outs.

As expected, $\texttt{Opt}(0)$ outperforms other methods in tracking $\underline{M}_4$ since it does have access to full knowledge and control. The main drawback is that the scalability of communication is low which makes the implementation difficult for large fleets. As expected, the DQoS index is the lowest for the optimization-based method since this method leads to more frequent cycling. By adding a lock in the optimization-based method it can be seen that the number of cyclings significantly drops, but this will worsen tracking, and PJM scores considerably. In fact, lock-out represents a plant model change and the results show that the optimization method is not adapting well to this change in the plant model (i.e., when reality hits). By comparing $\texttt{Fit}(15,3,\text{E})$ and $\texttt{Fit}(15,0,\text{NE})$ it can be seen that removing the cycling constraint and using Eq. (6) improves the tracking ability both in terms of PJM composite score and tracking. The same impact is seen for the 3-minute fitness update time. In the PEM method, we can see that randomizing packets improves the tracking significantly while decreasing the CQoS since there exist longer packets in this case which are more likely to exceed the temperature limit. The results show that the fitness-based method and PEM can provide tracking scores close to the optimization-based method while providing much higher M_6 and M_7 scores.

V. Conclusion

In this paper, a simulation-based analysis is used to validate a novel methodology for quantifying and comparing the performance of different DER coordination schemes. The methodology enables comparison across seven different, practical, and relevant metrics. For example, it is clearly shown that while the optimization-based method provides excellent tracking capability and CQoS, its high communication burden makes it

TABLE I: Comparing the performance of DER coordination schemes in terms of average and worst values of $X_{j,i}$ across $i=1,...,N_s$

Method	DQoS		CQoS		SoComm(bps)		Tracking (kW)		PJM score		SoComp(sec)		Availability	
	avg	worst	avg	worst	avg	worst	avg	worst	avg	worst	avg	worst	avg	worst
Fit(3,3,NE)	1608	2145	62.0	104	1.470	2.942	7.9	30.2	0.940	0.896	0.031	0.039	0.967	0.963
Fit(3,3,E)	1606	2139	63.4	101	1.470	2.942	7.5	28.7	0.940	0.896	0.137	0.148	0.967	0.964
Fit(3,0,E)	1619	2221	62.8	99	1.470	2.942	1.5	1.6	0.943	0.902	0.142	0.155	0.999	0.998
Fit(15,3,NE)	1774	2385	150.7	223	0.2978	0.2978	11.3	48.8	0.939	0.891	0.032	0.038	0.969	0.965
Fit(15,3,E)	1605	2145	62.5	103	0.2978	0.2978	7.5	30.2	0.940	0.896	0.137	0.151	0.967	0.963
Fit(15,0,E)	1617	2213	63.6	106	0.2978	0.2978	1.5	1.6	0.943	0.902	0.140	0.159	0.999	0.998
Opt(0)	49862	55449	0.063	1	68.11	68.12	4.6	6.6	0.944	0.905	9.117	10.98	0.999	0.999
Opt(3)	3558	4039	163	248	68.01	68.01	166.1	248.4	0.747	0.639	3.343	3.876	0.966	0.964
PEM(3,3)	3797	4188	100.6	151	0.107	0.116	65.1	108.5	0.855	0.739	0.021	0.055	0.995	0.992
PEM([1,5],3)	3665	4057	112.7	150	0.107	0.113	37.1	72.9	0.919	0.863	0.023	0.047	0.997	0.996

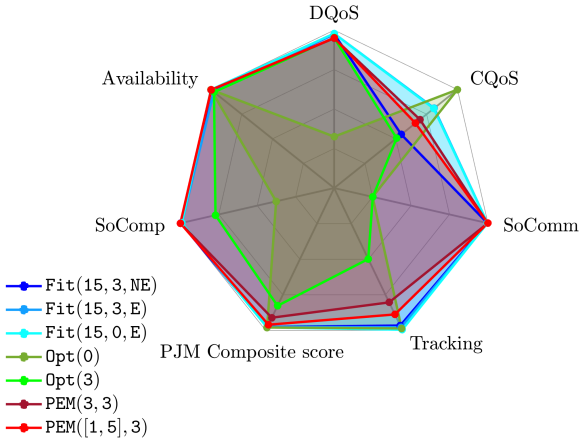

Fig. 2: Comparing mean performance of coordination schemes.

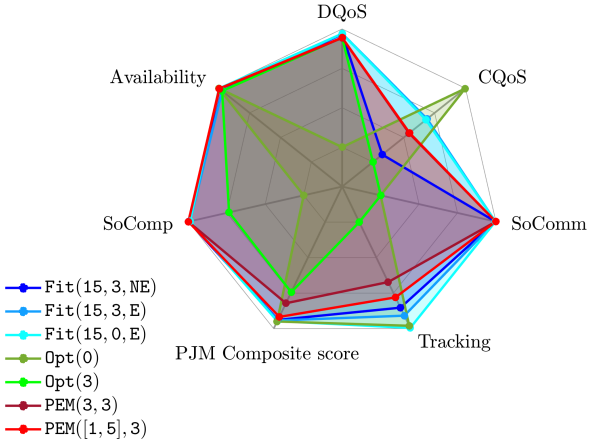

Fig. 3: Comparing worst performance of coordination schemes.

inappropriate for large-scale applications. It was also shown that by appropriate selection of the parameters, the proposed modified fitness-based method and PEM are able to provide comparable tracking capability to the optimization-based method with much lower communication requirements. Future work will focus on developing a comprehensive metric that quantifies the "efficiency of coordination" as well as considering network constraints/bottlenecks, and different types of devices/loads and grid services.

REFERENCES

[1] J. HU, G. YANG, K. KOK, Y. XUE, and H. W. BINDNER, "Transactive control: a framework for operating power systems characterized by high penetration of distributed energy resources," *Journal of Modern Power Systems and Clean Energy*, vol. 5, pp. 451–464, 2017.

[2] S. Brahma, A. Khurram, H. Ossareh, and M. Almassalkhi, "Optimal frequency regulation using packetized energy management," *IEEE Transactions on Smart Grid*, pp. 1–1, 2022.

[3] H. Mavalizadeh, L. A. Duffaut Espinosa, and M. R. Almassalkhi, "Decentralized frequency control using packet-based energy coordination," in *2020 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*, 2020, pp. 1–7.

[4] F. Moret and P. Pinson, "Energy collectives: A community and fairness based approach to future electricity markets," *IEEE Transactions on Power Systems*, vol. 34, no. 5, pp. 3994–4004, 2019.

[5] M. Zholbaryssov, C. N. Hadjicostis, and A. D. Dominguez-Garcia, "Privacy-preserving distributed coordination of distributed energy resources," in *2020 59th IEEE Conference on Decision and Control (CDC)*, 2020, pp. 4689–4696.

[6] J. Qi, Y. Kim, C. Chen, X. Lu, and J. Wang, "Demand response and smart buildings: A survey of control, communication, and cyber-physical security," *ACM Trans. Cyber-Phys. Syst.*, vol. 1, no. 4, oct 2017.

[7] D. Sloot, N. Lehmann, and A. Ardone, "Explaining and promoting participation in demand response programs: The role of rational and moral motivations among german energy consumers," *Energy Research & Social Science*, vol. 84, p. 102431, 2022.

[8] M. R. Sarker, M. A. Ortega-Vazquez, and D. S. Kirschen, "Optimal coordination and scheduling of demand response via monetary incentives," *IEEE Transactions on Smart Grid*, vol. 6, no. 3, pp. 1341–1352, 2015.

[9] F. Charbonnier, T. Morstyn, and M. D. McCulloch, "Coordination of resources at the edge of the electricity grid: Systematic review and taxonomy," *Applied Energy*, vol. 318, p. 119188, 2022.

[10] M. Almassalkhi, L. A. Duffaut Espinosa, P. D. H. Hines, J. Frolik, S. Paudyal, and M. Amini, *Asynchronous Coordination of Distributed Energy Resources with Packetized Energy Management*. New York, NY: Springer New York, 2018, pp. 333–361.

[11] A. Bhattacharya, J. Hansen, K. Kalsi, J. Lian, S. P. Nandanoori, H. Reeve, V. Adetola, F. Lin, T. Leichtman, S. N. Gourisetti, W. Hofer, S. Kundu, L. Marinovici, S. Niddodi, D. Vrabie, M. Chiodo, S. Yuan, and D. Wright, "Incentive-based control and coordination of distributed energy resources," PNNL, Tech. Rep. DOE-PNNL-28724, 05 2019.

[12] S. P. Nandanoori, I. Chakraborty, T. Ramachandran, and S. Kundu, "Identification and validation of virtual battery model for heterogeneous devices," in *2019 IEEE Power & Energy Society General Meeting (PESGM)*, 2019, pp. 1–5.

[13] A. Khurram, M. Amini, L. A. D. Espinosa, P. D. H. Hines, and M. R. Almassalkhi, "Real-time grid and der co-simulation platform for testing large-scale der coordination schemes," *IEEE Transactions on Smart Grid*, vol. 13, no. 6, pp. 4367–4378, 2022.

[14] M. Almassalkhi, J. Frolik, and P. Hines, "Packetizing the power grid: The rules of the internet can also balance electricity supply and demand," *IEEE Spectrum*, vol. 59, no. 2, pp. 42–47, 2022.

[15] S. P. Nandanoori, S. Kundu, D. Vrabie, K. Kalsi, and J. Lian, "Prioritized threshold allocation for distributed frequency response," in *2018 IEEE Conference on Control Technology and Applications (CCTA)*, 2018, pp. 237–244.

[16] O. Oyefeso, G. Ledva, M. Almassalkhi, I. Hiskens, and J. Mathieu, "Control of aggregate air-conditioning load using packetized energy concepts," in *2022 IEEE Conference on Control Technology and Applications (CCTA)*, 2022.

[17] S. Lloyd, "Least squares quantization in pcm," *IEEE Transactions on Information Theory*, vol. 28, no. 2, pp. 129–137, 1982.