Experimental Validation of an Explicit Power-Flow Primary Control in Microgrids

Lorenzo Reyes-Chamorro, *Member, IEEE,* Andrey Bernstein, *Member, IEEE,* Niek J. Bouman,
Enrica Scolari, *Student Member, IEEE,* Andreas M. Kettner, *Member, IEEE,* Benoit Cathiard,
Jean-Yves Le Boudec, *Fellow, IEEE* and Mario Paolone, *Senior Member, IEEE*

Abstract—The existing approaches to control electrical grids combine frequency and voltage controls at different time-scales. When applied in microgrids with stochastic distributed generation, grid quality of service problems may occur, such as under- or overvoltages as well as congestion of lines and transformers. The COMMELEC framework proposes to solve this compelling issue by performing explicit control of power flows with two novel strategies: (1) a common abstract model is used by resources to advertise their state in real time to a grid agent; (2) subsystems can be aggregated into virtual devices that hide their internal complexity in order to ensure scalability. While the framework has already been published in the literature, in this paper we present the first experimental validation of a practicable explicit power-flow primary control applied in a real-scale test-bed microgrid. We demonstrate how an explicit power-flow control solves the active and reactive power sharing problem in real time, easily allowing the microgrid to be dispatchable in real time (i.e. it is able to participate in energy markets) and capable of providing frequency support, while always maintaining quality of service.

Index Terms—Explicit Power-Flow Control, Real-Time Control, Real-Time Demand-Response, Primary Control, Microgrids.

I. INTRODUCTION

THE grand challenge of massively integrating volatile distributed generation into the power systems is strictly related to the evolution of their operation and control. The literature of the last decade has suggested two models for their future development [1]: (i) the *supergrid*, based on enhanced continental/intercontinental interconnections (mainly DC) [2]; or (ii) the *microgrid*, small medium/low voltage networks interfacing heterogeneous resources such as local generation, energy storage, and active customers, intelligently managed so that they are operated as independent cells capable of providing different services and operate as islands [3].

Irrespective of the model that will emerge, the control of heterogeneous distributed resources represents a fundamental challenge for both. This requires the definition of scalable and composable control methods that guarantee the optimal and feasible operation of distribution grids in order to satisfy local objectives (e.g., distribution grid quality of supply) as well as the provision of ancillary services to the external bulk transmission (e.g., primary and secondary frequency supports).

Several control methodologies have been proposed to achieve these goals (e.g. [4]), and the majority have been inspired by the time-layered approach traditionally adopted in power transmission systems, i.e. primary, secondary and tertiary controls, ranging from sub-second to hours time-scales respectively. In the context of microgrids, these three levels of control can be associated with a decision process that can be centralized (i.e., a dedicated central controller decides on the operation of the system resources) and/or decentralized (each element decides based on its own rules). In the current literature, the former is used for long-term, and the latter for short-term decisions. In particular, primary controls are deployed through fully decentralized schemes mainly relying on the use of droop control (e.g. [5]).

In this paper, we focus on the control problem associated to the time-scales of the primary and secondary controls. Moreover, we focus on microgrids, which are the smallest subsystems where such a control framework can be deployed, due to their small (if not null) inertia, and to the presence of heterogeneous and stochastic devices (e.g., distributed generation and storage). Specifically, we use the recently proposed COMMELEC control framework [6] (Composable Method for Real-Time Control of Active Distribution Networks with Explicit Power Setpoints), which is an *explicit power-flow primary control*. This means that it explicitly computes optimal nodal power injections/absorptions in order to achieve a global objective in real-time. This approach is in contrast to the use of local (e.g. proportional) controls (e.g., droop-based control strategies). It allows to steer an entire network as an equivalent energy resource, thus enabling an entire system to support the main grid by exploiting the flexibility of its components in real-time. In this paper, we report the first real-scale implementation of COMMELEC, thus demonstrating the flexibility of explicit power-flow control in microgrids.

The contributions of the paper are the following. Firstly, discuss the requirements for a primary power-flows control to be deployed in real microgrids (Section III), i.e., system awareness infrastructure, dedicated hardware for distributed resources and communication infrastructure. Secondly, describe the modifications that are needed to the internal structure of the COMMELEC grid agent solver in order to be applicable to a real system (Section IV). Thirdly, illustrate the application of the COMMELEC framework to control a set of heterogeneous energy resources (Section V). Finally, report on the first real-scale experiment that proves the applicability of an explicit power-flows control mechanism in microgrids (Section VI).

II. STATE OF THE ART

In the context of microgrids, the basic control functions have been adopted from those in bulk power systems. This adaption is usually different depending on the mode of operation: grid-connected or islanded. Due to the size of microgrids, the control functions can be combined hence simplifying the control architecture. For example, it is common to find strategies that solve both secondary and tertiary controls. We next present a summary of the state-of-the-art in microgrid controls.

A. Primary Controls

Primary controls[1] are mainly responsible for immediately responding to power imbalances in the grid. This becomes more difficult in microgrids due to the lack of inertia and the volatility of renewable distributed generators [7]. In particular, microgrids are typically characterized by having all (or almost all) the resources interfaced through power electronics [8]. Moreover, the resources usually have similar power ratings [9] (i.e. there is no *main* generator acting as slack).

Voltage-source inverters (VSIs) are typically used in micro-grids since they do not need an external reference to stay synchronized and can provide services like ride-through capability and power quality enhancement [10]. Droop controllers are commonly used, relying on the decoupling between the influence of active power P and reactive power Q on frequency f and voltage magnitude V, respectively [11]. The idea of droop control is to use proportional controllers:

$$f = f^* - m_f(P - P^*), \quad V = V^* - m_V(Q - Q^*), \quad (1)$$

where P^* and Q^* are the references of these controllers, f^* and V^* the expected values at the reference powers, and m_f and m_V the proportional gains.

B. Secondary Controls

As in traditional grids, secondary controls in microgrids aim at minimizing the frequency and voltage deviations with respect to their rated values. This is done by centrally modifying V^*, f^*, m_V and m_f of Eq. (1) according to the state of the grid and to the predefined operational bounds of the controllable distributed resources.

C. Tertiary Controls and Long-Term Energy Balance

The traditional tertiary control, as used in the bulk transmission network, does not play an important role in microgrids since secondary reserve does not typically exist in microgrids operation [12]. Instead, it is strictly linked to the long-term objectives, and it is typically referred to the actions of the so-called Energy Management System (EMS). Its main objective is to plan the lowest-cost operation of the microgrid or, in islanded operation, to guarantee the long-term availability of supply. Two categories can be found in literature: the fully centralized EMS and the distributed EMS.

D. Distributed and Decentralized Approaches

In this category, droop-based methods are well spread in implemented microgrids. However, the conventional method has disadvantages, among which the main ones are [13]:

- Lack of knowledge of the state of the energy storage, leading to non-secure and sub-optimal control decisions.
- Poor performance in grids with low X/R ratio of line longitudinal parameters.
- Instability issues when choosing small droop parameters (desirable for keeping the deviations small).
- Disregard of the uncertain resources dynamics, that can lead to limit violations and instability.
- Poor power sharing due to output impedance uncertainties that are result of the internal control laws.

In the following, we briefly discuss how the literature has proposed to tackle these issues.

1) Conventional Droop-based Methods: For low-voltage grids, the X/R ratio decreases. Therefore, the active and reactive power are both related with voltage and frequency deviations. A first solution is presented in [14]. It accounts for the full coupling between control and controlled variables. The control law is kept almost the same:

$$f = f^* - m_f(P' - P^*), \quad V = V^* - m_V(Q' - Q^*), \quad (2)$$

where P' and Q' are computed as

$$\begin{bmatrix} P' \\ Q' \end{bmatrix} = \begin{bmatrix} X/Z & -R/Z \\ R/Z & X/Z \end{bmatrix} \begin{bmatrix} P \\ Q \end{bmatrix}, \quad (3)$$

being P' and Q' the *modified* active and reactive powers.

Additionally, as the equivalent impedance seen by a converter is affected both by passive elements and its internal control laws, a further improvement is to add an outer loop known as *virtual impedance* [15]. It modifies the voltage reference of the internal control loop to follow a desired behavior [16], represented in the Laplace domain as $V_{ref}(s) = V(s) - Z_v(s)I(s)$, where V is the voltage given by the droop control, Z_v the virtual impedance, I the output current and V_{ref} the modified reference voltage.

2) Advanced Droop-based Methods: In order to cope with the instability issues and the lack of knowledge of the behavior of the uncertainty and the state of the storage devices, improved droop-based methods have been proposed in the literature.

As for stability, a modified droop-strategy (proportional-derivative) is proposed in [17], for resistive output impedance:

$$f = f^* - m_f(P' - P^*) - \hat{m}_f \frac{dP'}{dt}, \quad (4)$$

with $\hat{m}_f$ the gain for the derivative term and similarly for Q.

When the controlled resources are storage systems, the power flexibility that they can offer changes as a function of their state-of-charge (SoC). Indeed, when reaching the storage bounds (full charge or depletion), the flexibility is strictly limited to half of the power range, so the droop parameters have to be saturated. To avoid saturation and allow a smooth computation of the droop parameters, SoC-based droop strategies have been proposed (e.g. [18]).

---

[1]Also referred to as local control or real-time control.

3) Droop-Free Methods: In this category, decentralized multi-agent systems (MAS) are mainly found, where agents communicate either with each other or with a central controller. For example, a wireless-based multi-agent coordination method is proposed in [19], and a self-organizing architecture method in [20]. Both perform control by using consensus theory.

E. Fully Centralized Approaches

In this category, the literature tackles the problem also by using multi-agent methods. Here, we present only those that apply to primary controls. Reference [21] comprehensively analyses the way such systems are typically used in power systems. In this context, MAS are proposed as a step towards the distribution of control. Optimization goals in previously proposed methods (e.g. [22]) consider the operational costs of the system without accounting for the operational constraints of the grid. In particular, it relies on the availability of droop control that is suitably adjusted by MAS negotiations. However, this method neither expresses the specific state of the resources nor considers the grid state.

Recently, a *top-down* unified framework for primary and secondary frequency controls and congestion management has been presented in [23]. It distributes an overall optimization problem into the various resources for primary control, and computes a centralized decision for secondary control. Indeed, in this approach the problem distribution strictly depends on the nature of the problem itself and, therefore, changing the objective will require a re-engineering of the distribution of the control, which hinders its scalability.

F. The COMMELEC Framework

The COMMELEC framework [6] can be categorized as a droop-free distributed hierarchical-agents-based method. It uses a hierarchy of software *agents* to control a power grid. Each resource is equipped with a *resource agent* (RA), which is tailored to the technology peculiarities of the resource. A RA translates the internal state and constraints of its resource into an abstract, device-independent format, thereby hiding its internal details. The resulting *advertisement* message is sent to a *grid agent* (GA), which controls a subsystem consisting of a number of resources and the electrical grid interconnecting them.
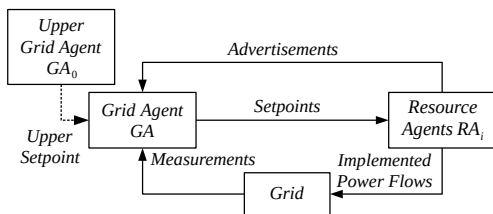

Fig. 1. Decision process in the COMMELEC framework.

An advertisement message consists of three elements.

Virtual cost and PQ profile. The *virtual cost* function, $\mathcal{C} : \mathcal{A} \to \mathbb{R}$, defines the cost to this resource of applying a setpoint[2] (P, Q) (namely, active and reactive power). Its domain of definition (a subset of $\mathbb{R}^2$) is called the PQ *profile* $\mathcal{A}$; it is the

set of setpoints that the resource is able to deploy according to its time-dependent constraints. The purpose of $\mathcal{C}$ is to quantify the preference of the resource for some setpoints and is not necessarily related to monetary value.

Belief function. The *belief function*, $\mathcal{B} : \mathcal{A} \to 2^{\mathbb{R}^2}$, is a *set-valued* function that quantifies the *uncertainty* in the operation of this resource. Assume that a RA receives from its GA a request to implement a given setpoint $(P, Q) \in \mathcal{A}$; the setpoint (P', Q') that is actually implemented, may differ from it, but must lie in its *belief*, $(P', Q') \in \mathcal{B}(P, Q)$. Highly controllable resources, such as batteries, are expected to have (almost) ideal beliefs, namely $\mathcal{B}(P, Q) = \{(P, Q)\}$. For uncertain resources such as PV, wind farms, or loads, the belief function will return larger sets to account for their volatility.

Since the advertisements are in a device-independent format, the GA software is generic and performs the same control for any grid (see Fig. 1). At every time step, it receives: (i) the advertisements from its RAs, (ii) the power setpoint requested from an upper GA and (iii) the estimation of the electrical state of the grid (using real-time methods, e.g., [24]). The GA steers the electrical state of its grid by explicitly setting power setpoints so that (i) the RAs virtual costs are minimized, (ii) the power setpoint requested by an upper GA is met and that (iii) the grid is in a *feasible state of operation*[3]. The process is repeated on a sub-second time scale, which is fast enough to cope with the fastest possible variations of distributed resources and slow enough to estimate the electrical state of the grid.

With respect to the discussed literature review, the COMM-ELEC framework is characterized by:

1) An abstract representation of resources and subsystems, a key element for a modular design.
2) Composability, i.e., entire subsystems can be abstracted in the same way as resources. This characteristic makes the approach scalable from low-voltage microgrids to medium-voltage distribution networks.
3) Real-time capability. Specifically, it proposes a formal approach to close the agent negotiation and deployment of setpoints in a sub-second time-scale.
4) Elimination of the constraint that distributed approaches cannot be adopted for primary control.
5) Optimal decoupling of the real-time control and the dispatch problems.

III. Deployment Requirements

In order to implement the decision process shown in Fig. 1, three elements are needed: (i) time-deterministic metering infrastructure, (ii) real-time sensing and processing hardware on each resource, and (iii) suitable communication network. We provide details about these elements, emphasizing the reasons that make them essential for the deployment of COMMELEC.

A. Microgrid metering system

In the COMMELEC framework, the GA needs to know the grid state at each computation-step so that a new set of power setpoints ensures safe and optimal operation of the grid.

---

[2]Negative power means consumption.

[3]Which refers to *static* (rather than dynamic) feasibility, defined in terms of the nodal voltage magnitudes and line currents
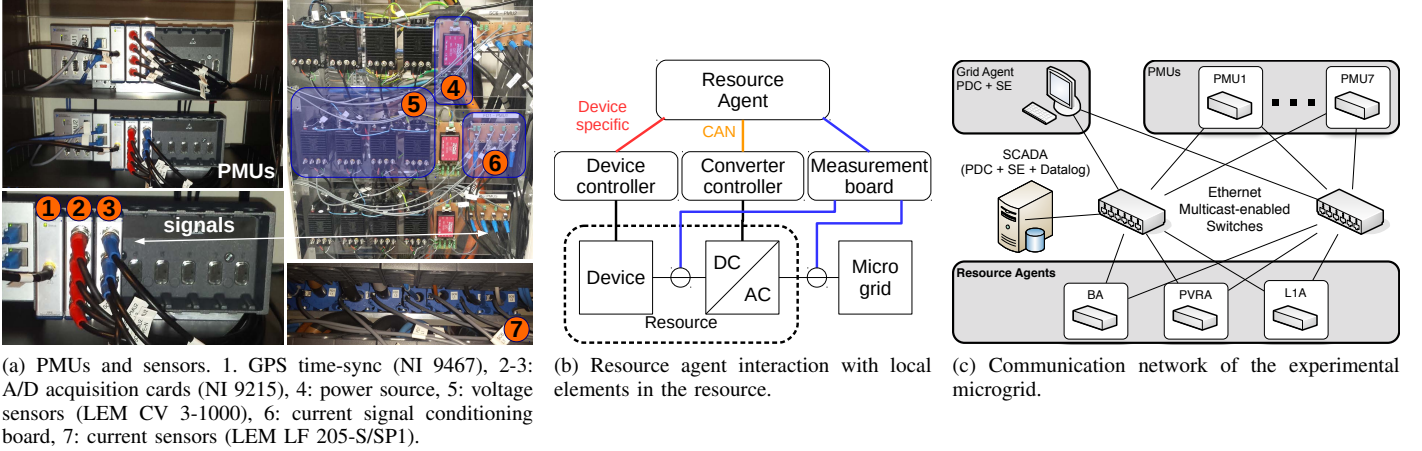
(a) PMUs and sensors. 1. GPS time-sync (NI 9467), 2-3: A/D acquisition cards (NI 9215), 4: power source, 5: voltage sensors (LEM CV 3-1000), 6: current signal conditioning board, 7: current sensors (LEM LF 205-S/SP1).

(b) Resource agent interaction with local elements in the resource.

(c) Communication network of the experimental microgrid.

Fig. 2. Metering system, resource agent hardware and communication network.

In this respect, we make use of Phasor Measurement Units (PMUs), since they can provide synchronized and time-tagged measurements, from which we can get a reliable grid state knowledge. Specifically, we use a PMU that satisfies every requirement defined in the standard for class-P PMUs [25].

The metering process (see Fig. 2a) starts with the voltage and current sensors, which feed the PMUs with scaled waveforms in order to estimate the synchrophasors. PMUs encapsulate UDP datagrams according to IEEE Std. C37.118.2-2011. A dedicated low-latency Phasor Data Concentrator is used [26]. It takes care of decapsulation of the PMU datagrams, and relies on a timeout-based buffer that uses the synchrophasors timestamps for time alignment. Moreover, it replaces missing measurements in order to feed applications with a consistent and complete set of data. This solution always ensures that the available measurements are forwarded in an acceptable time range (20ms in the case of PMUs streaming at 50fps) thus increases the determinism of the process. We use a discrete-Kalman-filter-based state estimator, which is suitable for 3-phase power systems and relies only on nodal synchrophasor measurements [27].

B. Hardware for Resource Agents

Every resource in the microgrid is equipped with a RA deployed in a NI CRIO 9068, which (i) verifies the feasibility of a requested power setpoint, (ii) commands the actual actuator in the resource to implement the verified power setpoint and (iii) provides an internal view of the resource in a generalized format in real-time. For achieving these tasks, a RA internally uses device-specific communication protocols to speak with both (device and converter) controllers and a measurement board from which it gets analog measurements (see Fig. 2b). In our case, variables subject to slow dynamics, such as the temperature, are sensed by the controllers, while electrical variables are directly measured.

C. Communication Infrastructure & API for Resource Agents

The microgrid is equipped with a dedicated IPv4 communication network (see Fig. 2c), that is redundant (using the iPRP redundancy protocol [28]) and secure (using Multicast security with elliptic curve authentication [29]). PMUs and RAs send their data using IP multicast to both the GA and a SCADA.

One key feature of the COMMELEC framework is its *device-independent* message format that enables resources to delegate decisions related to their control actions to a GA. In order to easily design a COMMELEC-compliant RA, we have developed a simple Application Programming Interface (API) consisting of two parts: an easy-to-use High-Level API, which supports a pre-defined set of commonly used resources and a Low-Level API that provides full access to the underlying device-independent message format.

Concretely, the High-Level API is provided by a *daemon* (background process) that runs on the RA machine as a "middleman" between the RA and the GA. The main benefits of this approach are that (i) a RA designer neither requires knowledge of the COMMELEC device-independent message format nor of the transport protocol and (ii) the interface between the RA and daemon, which is *device-dependent*, becomes very simple (only a few floating-point-typed parameters need to be exchanged). Hence, the daemon takes care of translating these device-dependent parameters into the device-independent message format used to communicate with the GA. This device-independent message format is defined on top of the *Cap'n Proto* serialization framework, which has very low parsing overhead. This is important at the GA side, which has to process advertisements from a potentially large number of RAs.

IV. THE GRID AGENT

The COMMELEC decision process, as seen in Fig. 1, relies on a robust multi-objective optimization procedure. Upon reception of the resource agent advertisements, the grid agent request will be chosen to minimize

$$F(u) = \sum_{i=1}^{n} \omega_i \tilde{\mathcal{C}}_i(P_i, Q_i) + \omega_0 \mathcal{C}_0(P_0(u), Q_0(u)) + J(u), \quad (5)$$

over the *admissible* power setpoint $u = (P_1, Q_1, ..., P_n, Q_n)$.

The function to be minimized is the weighted (ω) sum of the *normalized* virtual costs $\tilde{\mathcal{C}}_i$ of the resources (see Section IV-A for details), a penalty term $\mathcal{C}_0$ on the power flow (P_0, Q_0) at the point of connection with the upper-level grid, and a penalty term J on the nodal voltages $\{V_k\}$ and line currents $\{I_l\}$. The costs of the resources are convex by design. Moreover, the two penalty terms, as functions of $(P_0(u), Q_0(u))$

and $(\{V_k(u)\}, \{I_l(u)\})$, are designed convex and, at each COMMELEC iteration, we compute the voltage and current sensitivity coefficients [30] to replace $P_0(u), Q_0(u), \{V_k(u)\}$ and $\{I_l(u)\}$ by their first-order Taylor approximations. The resulting approximation is thus convex. The admissibility of a setpoint u is defined in terms of the uncertainty sets provided by the advertised belief functions as defined in [6].

The optimization procedure is based on a single iteration of the gradient projection algorithm per COMMELEC cycle. More precisely, the new setpoint selected by the GA is

$$u = \mathcal{P}[\hat{x} - \alpha \nabla_u F(\hat{x})], \qquad (6)$$

where $\mathcal{P}$ is the projection operator to the admissible space, $\hat{x}$ is an estimate of the currently implemented setpoint, and α is the gradient step-size.

Observe that, the GA optimization procedure is sensitive to the scale of the different terms in the objective function (5). In particular, the virtual cost functions $\mathcal{C}_i$ obtained from different RAs can be of different magnitude, and thus can affect in an arbitrary, and possibly unfair, way the gradient update (6). Thus, some normalization of the cost terms is required.

Also, since we are tracking the optimal solution of a time-varying optimization problem, correctly selecting the gradient step-size is essential. We want to prevent the GA output from oscillating around the optimal trajectory or from diverging from it (step-size parameter too big). It is also important to ensure that the convergence rate is not too slow (step-size parameter too small), especially in case of fast-changing objectives.

Next, we show how this implementation aspects are tackled. In our experimental setup the GA is written in C++, and currently runs on a Scientific Linux 7.2 64-bit desktop workstation. The GA code can be easily cross-compiled to other platforms, including embedded ones.

A. Normalization of Cost Functions

Inspired by a classical method for choosing the gradient step-size in convex (single objective) optimization theory, in this paper we normalize the costs using the Lipschitz constant of the cost gradient. Also, we use *convex quadratic* cost functions for the different RAs. Moreover, the cost functions are separable for P and Q. Hence, without loss of generality, we consider a quadratic function that depends only on P, of the form $\mathcal{C}(P) = aP^2 + bP, a > 0$. We normalize this function using the Lipschitz constant for $\mathcal{C}'$ (the derivative of $\mathcal{C}$) that is simply given by the second derivative, namely $L = 2a$. The normalized cost is then given by $\tilde{\mathcal{C}}(P) = P^2/2 - P^*P$, where $P^* \triangleq -b/2a$ is the stationary setpoint for the RA. Thus, the derivative of the normalized cost is *affine*: $\tilde{\mathcal{C}}'(P) = P - P^*$.

B. Selection of Gradient Step-Size

To avoid extreme behaviours given by too-small or too-big gradient step-size, we decide to select the step-size α at each iteration, in order to minimize $f(\alpha) = F(\hat{x} - \alpha \nabla_u F(\hat{x}))$. Since the overall cost F is convex, the aforementioned one-dimensional optimization problem is still convex and not computationally expensive. We define lower and upper bounds for the step-size and efficiently find a solution by binary search.
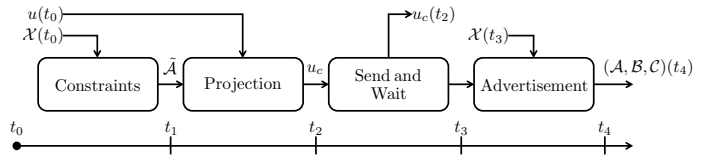

Fig. 3. Full process performed by a generic resource agent. It is expected that the computation time $t_4 - t_0 < 30$ms.

C. Enabling Frequency Support to the Main Grid

We enable the microgrid to provide frequency support by including the following penalty term into its overall objective function: $(\hat{P}_0[k] - P_{\text{target}}[k])^2$, where $\hat{P}_0[k]$ denotes the (estimated) active power flow at the connection point and $P_{\text{target}}[k] := \hat{P}_0[k] + \nu^{-1}(f_{\text{nom}} - f[k])$, where $\nu \geq 0$ is the frequency-droop parameter (unit: [Hz/W]), f_{nom} is the nominal frequency and $f[k]$ is the measured frequency at time k.

V. RESOURCE AGENTS DESIGN

A. Resource Agents Process

In general, a RA implements four consecutive tasks, as depicted in Figure 3.

1) Constraints: Upon receiving a request $u(t_0)$, it retrieves the resource state $\mathcal{X}(t_0)$ and computes the current PQ-capabilities $\tilde{\mathcal{A}}$. This set may differ from the previously computed $\mathcal{A}$ due to delays or innacuracies.

2) Projection: At t_1, it verifies the feasibility of $u(t_0)$ according to $\tilde{\mathcal{A}}$. That is, it performs a simplified Euclidean projection onto $\tilde{\mathcal{A}}$. The result is the command u_c. For simplification, we assume this set is the intersection $\tilde{\mathcal{A}} = \tilde{\mathcal{A}}^r \cap \tilde{\mathcal{A}}^p$, where $\tilde{\mathcal{A}}^r$ are linear constraints in rectangular and $\tilde{\mathcal{A}}^p$ in polar coordinates. Then, using $u_0^* = u(t_0)$, we alternately compute (changing the coordinates on each iteration):

$$u_{j+1}^* = \inf_{x \in \tilde{\mathcal{A}}^i, i=\{p,r\}} \{||x - u_j^*||\}, \qquad (7)$$

where $|| \cdot ||$ is the ℓ^2-norm. The iteration ends when $||u_{j+1}^* - u_j^*|| < \epsilon$, for $\epsilon > 0$ a small number and $u_c = u_{j+1}^*$.

3) Send and Wait: Using a resource-dependent protocol, the agent sends $u_c(t_2)$ to the actuator and waits until the setpoint is implemented. The waiting time depends on the resource's nature (e.g. on ramping constraints and controller parameters).

4) Advertisement: After this delay, it re-obtains the state $\mathcal{X}(t_3)$ when the setpoint has been deployed, and it computes the advertisement $(\mathcal{A}, \mathcal{B}, \mathcal{C})(t_4)$ to be sent to the GA [31].

B. Operation of a Resource Agent

For any RA, the sending/receiving message cycle is repeated continuously and endlessly during normal operation: a RA computes an advertisement just after receiving a request. The whole cycle of a RA is described next (see Figure 4).

1) At any timestep k, the RA receives a new request $(P, Q)[k]$. The request is to maintain this power fixed from k to $k + 1 = t_0 + \Delta t$, where Δt is the expected request rate. Observe that, $(P, Q)[k]$ has been computed (at the GA level) with the knowledge of $(\mathcal{A}, \mathcal{B}, \mathcal{C})[k - 1]$.
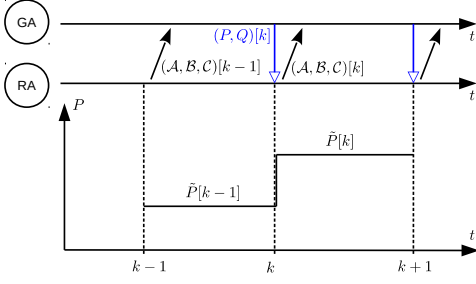2) Perform the Resource Agent Process (Section V-A).

Fig. 4. Representation of the messages between GA and RA in time. The lower diagram shows the implementation of the requested setpoint. For simplicity, only active power is shown.

3) Send the advertisement $(\mathcal{A}, \mathcal{B}, \mathcal{C})[k+1]$ to the GA. This advertisement should ensure that $(P, Q)[k]$ maintains the resource safe until timestep $k+1$.
4) Wait for a new request from GA.

In case a request takes even longer than Δt, the resource is expected to continue injecting/absorbing the same nodal power unless it violates an internal constraint. This delay may eventually cause a violation on the grid operational constraints (related to the term $J(u)$ in Eq. (5)) or to lose the track of the external setpoint (related to $\mathcal{C}_0$ in Eq. (5)). Therefore, as in our practical case a COMMELEC cycle takes around 100ms (see Section VI-F), we take $\Delta t = 300$ms. A value conservative enough to account for possible delays caused by external sources (communication, other RAs or the GA).

In the following, we describe the different RAs and how their advertisements are computed.

C. Resource Agent for Batteries

In our microgrid, the battery energy storage system has a 25kW/25kWh power/energy rating. It is based on Lithium Titanate cells and it is monitored by a local battery management system (BMS). The battery uses a 4-quadrants converter, which directly receives power setpoints. In what follows we describe how to abstract its internal state and constraints.

1) PQ Profile: The battery admissible set in the PQ plane is defined by the disk $\mathcal{A}^p = \{(\rho, \theta) | \rho \leq S_r = 25\text{kVA}\}$, and $\mathcal{A}^r = \{(P, Q) | \underline{P} \leq P \leq \overline{P}\}$, corresponding to the AC active power limits. $\underline{P}$ and $\overline{P}$ depend on the DC constraints, the converter maximum current, and its efficiency.

To find the limits of DC power at the next timestep, we use the two-time-constant (TTC) model in Fig. 5, whose parameters depend on the SoC. We assume that the SoC is fixed between two setpoint implementations and it is given by the BMS. The battery parameters have been assessed off-line as a function of the SoC as proposed in [32].
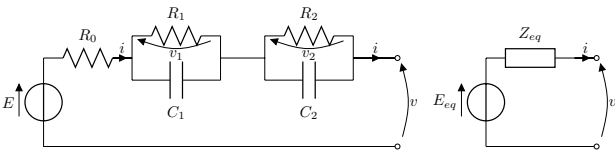

Fig. 5. TTC model (left) and Thevenin equivalent model (right).

The TTC model is formulated by using discrete stochastic state-space equations where k is the generic time-step, u is the

input vector, v the output, w and e are white noises associated with the process and the measurements, respectively.

$$x[k+1] = Ax[k] + Bu[k] + Gw[k], \qquad (8)$$
$$v[k] = Cx[k] + Du[k] + e[k]. \qquad (9)$$

A, B, C and D are functions of the model parameters, and therefore can change with the SoC, and G is the state disturbance matrix. In particular, for the TTC model we have:

$$A = \begin{bmatrix} -1/\tau_1 & 0 \\ 0 & -1/\tau_2 \end{bmatrix}, \quad B = \begin{bmatrix} 1/C_1 & 0 \\ 1/C_2 & 0 \end{bmatrix}, \quad x = \begin{bmatrix} v_1 \\ v_2 \end{bmatrix},$$

$$C = \begin{bmatrix} -1 & -1 \end{bmatrix}, \qquad D = \begin{bmatrix} -R_0 & E \end{bmatrix}, \quad u = \begin{bmatrix} i & 1 \end{bmatrix}^T,$$

with $\tau_i = R_i C_i$ for $i = 1, 2$. For this model, the residuals associated with the one-step-ahead prediction are not correlated, then we can apply Kalman filter for state estimation [33].

According to the battery specifications, we set: $\underline{v} = 442$V and $\overline{v} = 663$V. The converter has a maximum current of $\overline{i} = 41$A. To simplify the computation of the power limitations, we identify a simpler battery-circuit composed by its Thevenin equivalent (see Fig. 5 right). We define E_{eq} as the equivalent voltage and Z_{eq} as the equivalent passive element with $E_{eq}[k] = E[k] - v_1[k-1]e^{\frac{-\Delta T}{\tau_1[k]}} - v_2[k-1]e^{\frac{-\Delta T}{\tau_2[k]}}$ and $Z_{eq}[k] = R_0[k] + R_1[k](1 - e^{\frac{-\Delta T}{\tau_1[k]}}) + R_2[k](1 - e^{\frac{-\Delta T}{\tau_2[k]}})$, where ΔT is the RA sampling period (in principle $\Delta T = \Delta t$). The DC power p can be computed as a function of a voltage v_a as:

$$p(v_a) = \frac{(E_{eq} - v_a)v_a}{Z_{eq}}. \qquad (10)$$

Considering $\overline{u} = \max(\underline{v}, E_{eq} - Z_{eq}\overline{i}, E_{eq}/2)$ and $\underline{u} = \min(\overline{v}, E_{eq} + Z_{eq}\overline{i})$, and replacing v_a in (10), we can compute the maximum and minimum DC power, respectively. We then use a transfer function to compute AC active power from DC power that has been experimentally determined off-line.

2) Virtual Cost: We define a cost function that tracks a target SoC that represents an admissible internal state of the batteries given by a long-term scheduler [34].

3) Belief Function: We assume that this device has no uncertainty in the setpoint deployment, thus $\mathcal{B}$ is a singleton.

D. Resource Agent for Photovoltaic Plants

Our generator resource is an uncontrollable 20kW PV-plant connected to two 10kW three-phase inverters always working in maximum-power-point-tracking (MPPT) mode. The PV Agent (PVA) is hence meant to predict the uncertainty of the generated power.

1) PQ Profile: The PQ profile is simply the actual measured active and reactive power.

2) Virtual Cost: As the resource cannot be controlled, we set the advertised virtual cost equal to zero.

3) Belief Function: In order to quantify $\mathcal{B}$ and, with it, the stochasticity of the PV resource, we implement a real-time bound predictor using the algorithm described in [35]. The method is based on the experimentally observed correlations between the time derivative of the AC active power P_{pv} and the errors caused by a generic point forecast method. It allows having accurate prediction intervals while being computationally efficient. For the reactive power, as the system works in MPPT mode, we consider a small variation around the actual implemented value.
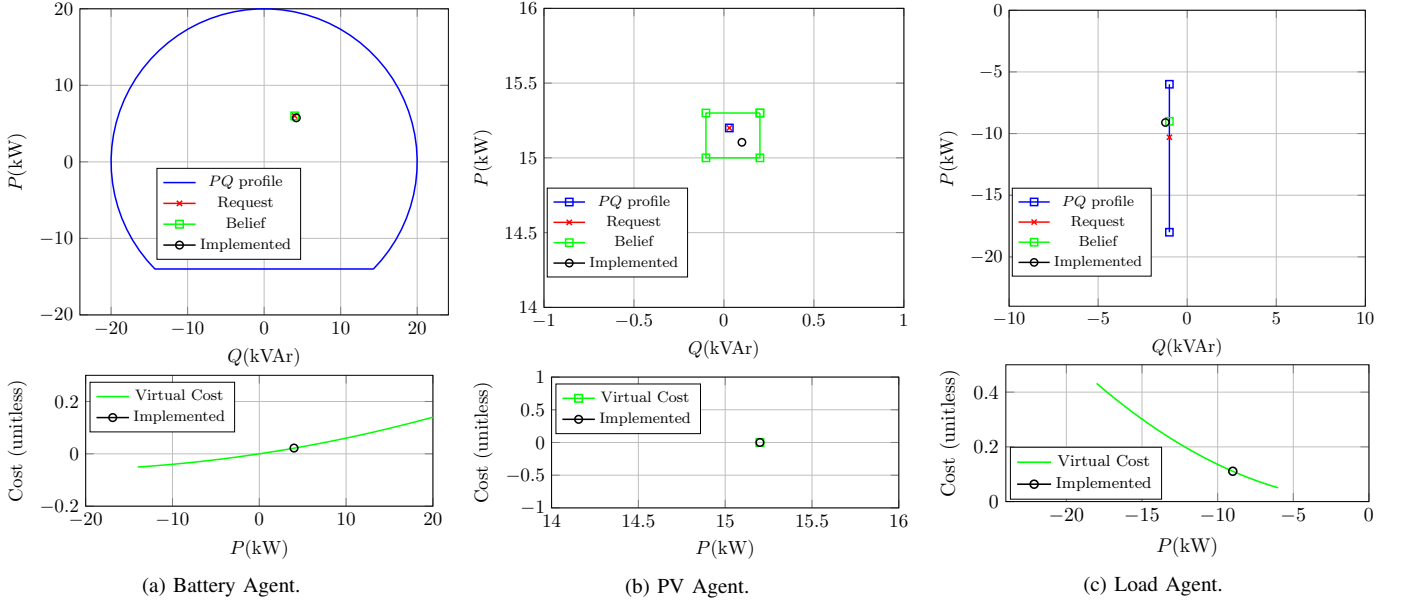
Fig. 6. Examples of Advertisements for every Resource Agent. Data taken from the SCADA.

E. Resource Agent for Electric-Heating Loads

We emulate the load using three controllable single-phase power converters. Indeed, a (P, Q) setpoint is translated into current amplitude and phase for the 3 phases.

Concretely, we emulate a 24kW residential building that has electrical space-heating. We model it with discrete state-space equations as described in [36]. It consists of 8 rooms, each of which is equipped with an electrical heater of 3kW (power factor 0.995) that can only be turned on or off. The controlled variable is the room temperature, which needs to remain within a given comfort zone $[\underline{T}, \overline{T}]$.
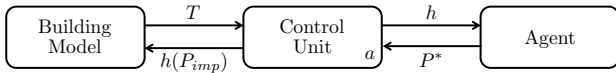
Fig. 7. Block diagram of the emulated load.

Fig. 7 shows the load emulation. In particular, the Control Unit (CU) gets the rooms temperatures T from the model and the command P_{imp} from the agent to compute the heating state h and the availability of each room a (a room can be locked to avoid constant switching). It dispatches the requested power to the available rooms, following a list where the rooms with the lower temperature are on top. The updated rooms state is then sent to the agent and to the building model. As the emulated load is discrete and the GA computes setpoints in a continuous way, the agent has to round the setpoint. This causes an error $\epsilon[k+1] = \epsilon[k] + k_p(P^*[k] - P_{imp}[k])$, that is reduced in average over time. The value of k_p determines the speed of the control action. Thus, at each timestep k, the CU sets $P_{imp} = \lceil (P^* + \epsilon)/m \rceil m$, where m is the specified increment (3kW).

1) PQ Profile: We always guarantee that the rooms' temperature are within the comfort bounds. Thus, we compute the feasible operation set as $\mathcal{A} = \{(P, Q) | \underline{P} \le P \le \overline{P}, Q = 0\}$. $\overline{P}$ is the sum of the power of the available heaters and whose room temperature is lower than $\underline{T}$. $\underline{P}$ is the maximum possible

power ($P_r = 24$kW) minus the consumption of rooms that are either not available or whose temperature is higher than $\overline{T}$.

2) Virtual Cost: We steer the rooms' temperature to the average temperature of the comfort interval, $\hat{T}$. The cost function is then defined as $\mathcal{C}(P) = aP^2 + bP$, with $a = 1/P_r^2$, $b = 2ea$, where e is the sum of the consumption of the rooms whose temperature is lower than $\hat{T}$.

3) Belief Function: In this case, the belief uses the error ϵ to indicate the expected implemented value for a given setpoint, namely $\mathcal{B}(P, Q) = (\lceil (P + \epsilon)/m \rceil m, 0)$. Note that, for every value of P, the result is a single point that is a function of ϵ.

VI. EXPERIMENTAL RESULTS

The results described in this section refer to the application of the COMMELEC framework to a real-scale microgrid, which is an implementation at realistic scale of the CIGRÉ low voltage (400V at 50Hz) microgrid benchmark defined in [37] (see Figure 8). For the sake of interpretation of the results, we decided to keep the size of the experiments to the minimum set of resources than can show the potential of the control framework. Thus, we consider a subset of the energy resources listed in [37]: a load (L1), a battery (B) and a photovoltaic plant (PVR) whose agents are defined as presented in Section V.

The microgrid of Fig. 8 is connected to a 20kV grid at bus B01 via a suitable transformer. The line that connects to the transformer, L01, has a current limit of 30A, i.e. a power transfer limit of ca. 21kVA, less than the maximum consumption of the loads. Such a scenario is typical for grid-tied microgrids where the investment made into local generation is partly compensated by a reduction of the peak power at the connection with the main grid.

We measured the latencies and found that state estimation accounts for the largest part of the delay (this is due to the need to sample up to three cycles of the AC waveform; at 50Hz this makes up to 60ms). In contrast, setpoint computation takes on average ≤ 2ms (See Section VI-F).

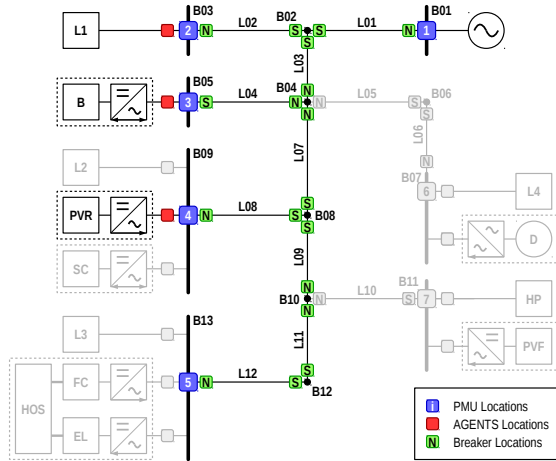Fig. 8. The experimental microgrid. All elements used in the experiments of this paper are shown solid print.

Fig. 9. Safe and Optimal Operation Experimental Results.

A. Safe and Optimal Operation of the Microgrid

In this experiment, we consider a daytime period where the objective of the battery is to be fully charged before the night. This may happen, for instance, if the cost of energy is low during the day (due to renewable penetration) and high during the night. A dispatcher would prioritize to store energy during the day and to sell during the night [38]. Another possibility is that the dispatcher foresees that the load during night exceeds the power that can be imported from the grid (e.g. due to the limited power transfer capability of line L01) and that the batteries will need to have enough energy to supply it (together with the grid) during the night. We also assume that, in the same considered time-window, the load has to consume (e.g. as the buildings are occupied). Therefore, the microgrid would like to store energy when the sun is available without disrupting the temperature comfort in the building. In this scenario, B and L1 are thus competing to have access to the electricity generated by PVR.

At $t = 0$s in Fig. 9, the battery charges close to its maximum allowed power (25kW) but, given the congestion in L01 and the fact that the temperature in the building is in the comfort zone, the consumption of L1 is reduced by the GA. Note that, given by the gradient of the penalty function associated to L01, the GA opts to keep an operational margin of its current. As the battery is charging at full power, its maximum consumption (defined by the PQ profile) is reduced in time, allowing the GA to slightly increment the consumption of L1, without violating the L01 current limit.

The congestion management at L01 is clear around $t = 400$s. L1 has to increase its consumption to 24kW, for example, as the temperature of several rooms is coming outside of the comfort zone. However, as the resulting current in line L01 is too close to its limit, the GA requests the battery to temporarily counteract by reducing its charge (increase the active power), and reducing its reactive power compensation. At $t = 405$s, the opposite takes place, the load needs less power and the battery starts gradually to consume more. We show, in particular, that COMMELEC exploits the flexibility of the building when it exists, but when the comfort zone is
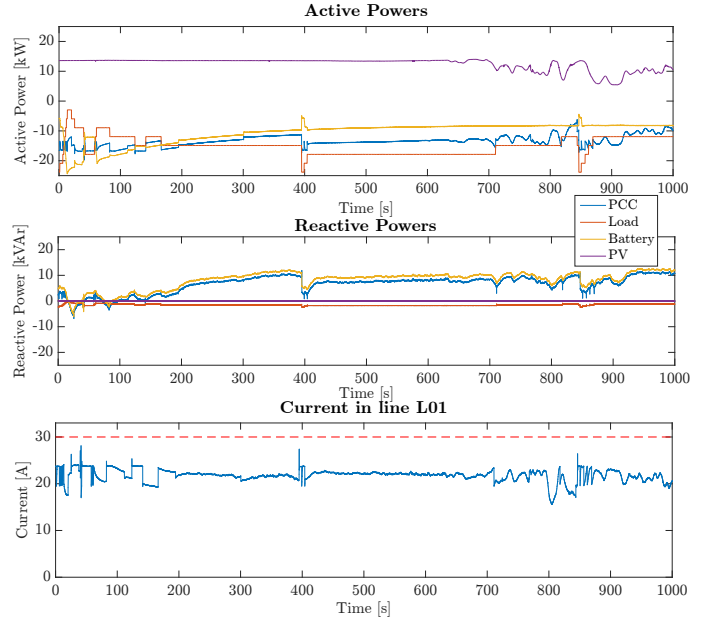
not respected, the building needs more power and, since the battery has less strict goals, the power load is increased. This experiment explicitly shows how the COMMELEC framework, despite being agnostic of the internal details of the resources, is able to perform real-time demand-response, accounting for the grid operational constraints.

B. Real-Time Dispatchability of the Microgrid

In this experiment, we show how the framework can be used to achieve the microgrid dispatchability. We assume that a dispatch plan is computed ahead of time. The goal of the real-time control is to follow the agreed-upon plan, as deviations are economically penalized [39]. The main feature of dispatchability is that coordination follows the commitment of operators on a predetermined dispatch plan. This is a simple yet reliable mechanism which reduces the complexity of the setup in terms of real-time communication and coordination requirements compared to, for example, real-time generation re-dispatch based on transactional control approaches. Thus, it is suitable to be implemented even in the current vertically operated power systems [33], [40].

Concretely, the microgrid, as a whole, needs to track a reference signal at B01. For this purpose, the GA simply activates a power-tracking function in its objective ($\mathcal{C}_0$ in Eq. (5)). Fig. 10 shows the results for tracking a predefined reference for active power and null for reactive power (in green). As one can see, the external setpoint is tracked with high fidelity, except in the cases when the current in line L01 is close to its ampacity. It is worth observing that this is achieved independent of the variability of the PV production. In the case of reactive power, as the load has a power factor of 0.995, its reactive power curve follows the active power one. Therefore, in order to follow the external request, the battery is requested to compensate.
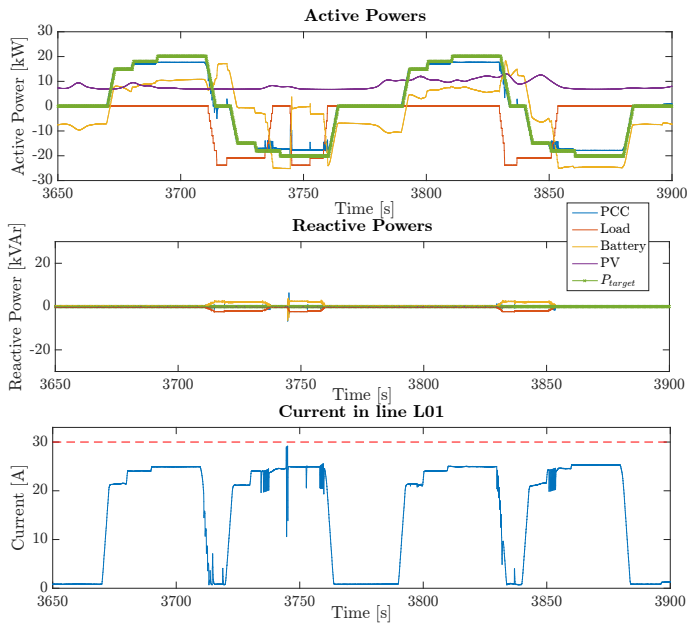
Fig. 10. Real-Time Dispatchability Experimental Results.

Fig. 11. Frequency Support Experimental Results.

C. Microgrid Provides Frequency Support to Main Grid

In this experiment, we show the capability of the microgrid to provide primary frequency support to the external grid. More precisely, assuming that the GA has the power-tracking function activated, as shown previously in Section VI-B, it can also enable a superposed frequency-support function. In other words, as participant in the energy regulation market, the microgrid can additionally deliver ancillary services to compensate the deviations of the frequency from its nominal power in the intra-second time-scale, namely, the microgrid, as a whole, provides primary frequency control to the main grid. The superposition of the power tracking mode and the frequency support is possible since the frequency deviations have a zero mean in the time scale of tens of seconds. Therefore, the two objectives can be reasonably decoupled. Concretely, the GA will implement a proportional compensation with respect to the grid frequency deviation with respect to the rated value of 50Hz (namely, the conventional droop-control). We have defined an a-priori microgrid regulating energy of 200kW/Hz.

We show in Fig. 11 the results for tracking a predefined dispatch plan (green line), while superposing the expected power at the connection point that compensates for the frequency deviations (light blue line). We see that, even with the volatility of the solar irradiance, the GA is able to steer the flexible resources to provide frequency support. Indeed, the power flow at B01 follows the shape of the frequency signal (mirrored vertically) by utilizing both the battery and the load[4]. Note that, when the current at line L01 gets close to the limit, there is a gap between the ideal frequency support and the actual realization, showing that the congestion management is always active. Given that the microgrid controller can follow the main-grid frequency variations, the microgrid could provide primary

reserve by participating in the ancillary services market and thus give frequency support to the main grid to which it is connected.

D. Autonomous Operation

In this experiment, we show the operation of the microgrid in islanded mode, being controlled by the COMMELEC framework. In order to perform this test, we have black-started the system using the battery as the slack resource, thus operating as a voltage-source-converter (grid-forming mode), while the rest of the resources work as before as current-source-converters (grid-feeding mode).

The main change in the GA, is the definition of the reference bus which, in islanded conditions, is not anymore the upper-level MV grid. The advertisements defined in Section V are used in the same way for all resources apart from the battery (i.e. the new slack resource), whose PQ profile is now used as a constraint for computing admissible power setpoints (for more details please refer to [6]).

The results in Fig. 12 show that the system is able to run flawlessly by following the safety and optimality objectives of Section VI-A. In this particular case, we have intentionally started the system in a situation where the battery needs to be discharged and the load needs to consume power.As the battery is the slack resource, it has priority to follow its goal (i.e. more weight in Eq. (5)). Hence, when it reaches its target SoC (set to 37% in this case) at around $t = 5600s$, the load is requested to decrease its consumption that has the flexibility to stop consuming. However, later at around $t = 6300s$, the load needs to consume again (to avoid the violation of the comfort bounds as indicated in Section V-E).

Note that, the voltage profile of all the nodes in the grid is very variable during the islanded operation, as expected for a system with negligible inertia. Yet, the voltage magnitude of all nodes is strictly kept in feasible conditions (see Fig. 12). In

---

[4]It is interesting to see that, because of the discrete action of the building emulator, some small spikes are present in the actual implemented power flow at B01. This behavior is a natural consequence of the discrete nature of the controller.
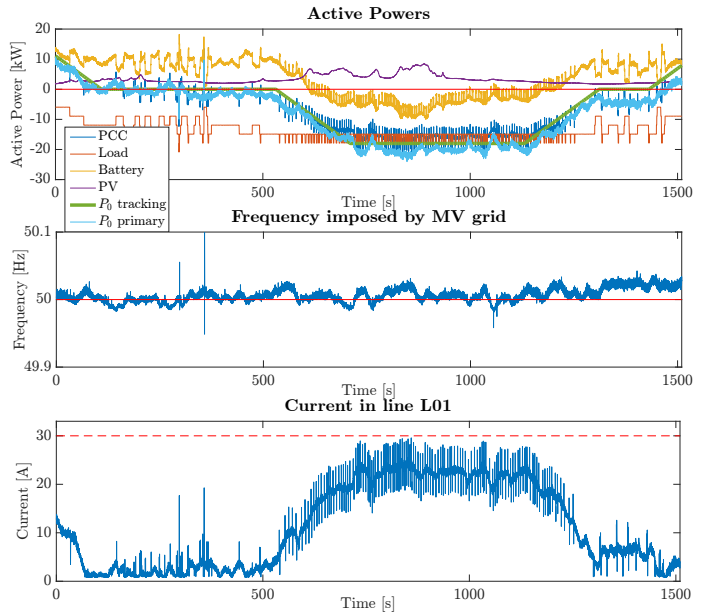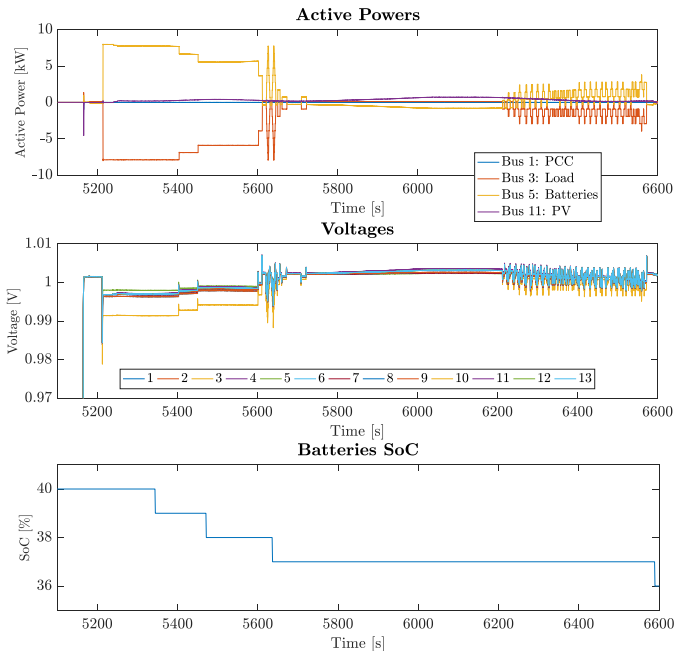
Fig. 12. Autonomous Operation Results.

this case, we see that the GA is still able to steer the resources to their own target operation while respecting the operational constraints of the grid. This analysis shows the flexibility of COMMELEC to control in a seamless way microgrids operating in generic conditions.

E. Communication Failure and Packet Losses

In this section, we show the impact of unreliable communication in our setup. As mentioned in Section III and illustrated in Figure 2c, we use the iPRP redundancy protocol in a communication system with two independent networks. When a packet loss or communication delay takes place in one path, this is repaired by the iPRP via the other path with zero-delay [28]. Note that, as the control loop is executed at a fast pace (in our setup every 100ms), the sporadic loss of packets will not largely affect the system behaviour unless there is a high power jump (this is accounted for in the belief function). Furthermore, as soon as the GA receives a subsequent packet, the loss does not have any further impact because the GA decisions are based on the last received advertisement and measurement.

For challenging our setup and making the impact of communication reliability more visible, we have produced a scenario where both communication paths (between the switch and GA) have a 5% rate of packets loss when the microgrid is operating in primary control on top of dispatch mode (as in Section VI-C). We show the impact on tracking the input signal for two cases in the presence of packet losses: (i) using only one communication path and (ii) using two communication paths and implementing iPRP (note that, the expected packet loss by using two paths and iPRP is reduced to 0.25%). Fig. 13 shows the behaviour of the system for the first case, where some sporadic losses are noticeable in the profile, while Fig. 14 illustrates the second case
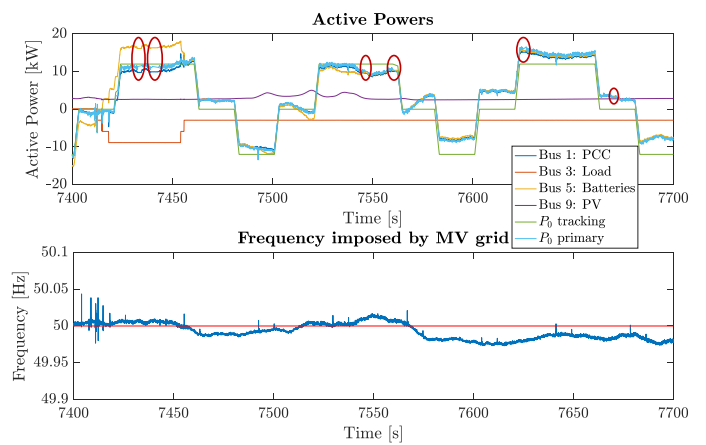

Fig. 13. Behaviour with Lossy Communication Link. The red circles show the impact of the losses.
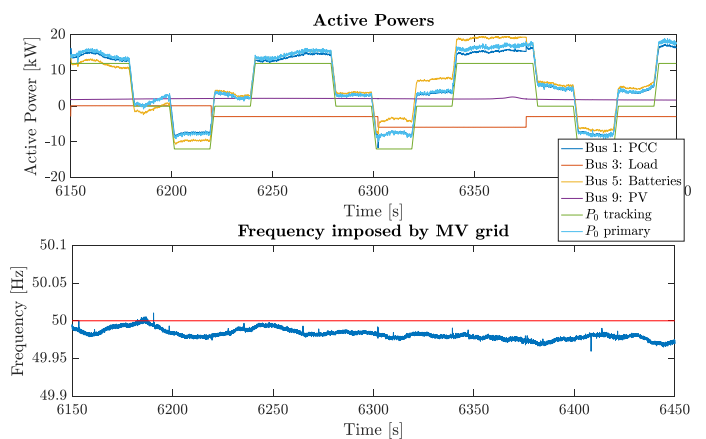

Fig. 14. Behaviour with Lossy Communication Links using iPRP.

where no losses are noticed[5]. As it can be seen, the adoption of the iPRP renders the control framework resilient against communication reliability issues.

F. Latency Assessment

In Fig. 15 we show the results of the experimental evaluation of the time-latencies incurred in the involved computation processes and their specific contribution to the overall COMMELEC cycle. The curves in Fig. 15 show the cumulative distribution of the latencies of each process. The beginning of the COMMELEC cycle, $t = 0$, is denoted by the instant when the GA sends a new set of setpoints to all RAs. After this, the GA waits until it receives the advertisements from the RAs. Each RA incurs, in general, in a different delay, which depends on its technology. Only after having received all advertisements, the GA waits until receiving an updated estimated state. Since the state estimation has to sample three cycles of the AC waveform to get a steady measurement (the PMU of [25] uses a Hanning window of 60ms), this stage represents the most important contribution to the overall delay. Only after getting the updated state, a new setpoint computation is issued, which

---

[5]Note that, we cannot reproduce the very same input in both cases due to the variable behaviour of the frequency imposed by the MV grid.

takes on average $\leq$ 2ms. A full COMMELEC cycle, as for the experimental validation here presented, takes less than 100ms.
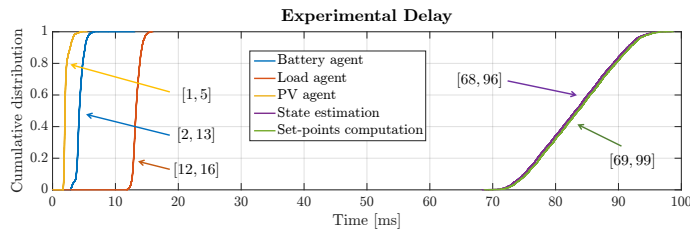

Fig. 15. Latency assessment. Cumulative distributions of the time latencies of all the processes measured from the instant when the GA setpoints are sent. The brackets represent the minimum initial time and maximum final time of each process.

VII. CONCLUSION

In this paper we have presented the experimental validation of the COMMELEC framework implemented in a real-scale microgrid deployed in a laboratory. We provide details on the deployment requirements for having access to real-time state estimation, fast reaction of resources for implementing setpoints and reliable communication. We show that the COMMELEC framework is able to optimally exploit the flexibility of heterogeneous resources without the need to access their internal states and parameters. The results also show that the framework is flexible and easily adaptable to cope with different goals such as microgrid dispatchability and frequency support, which can be superposed with no impact to the safety of the grid. This enables the microgrid to actively participate in the energy market and to provide ancillary services to the main grid in real-time. We demonstrated that the whole COMMELEC control cycle is executed within 100ms, thus guaranteeing its applicability to perform generic real-time control functions.

REFERENCES

[1] N. Hatziargyriou et al. CIGRE WG Network of the Future. Electricity supply systems of the future. *Electra*, 256:42–49, June 2011.

[2] S. Gordon. Supergrid to the rescue. *Power Engineer*, 20(5):30–33, 2006.

[3] N. Hatziargyriou, editor. *Microgrids: Architectures and Control*. John Wiley and Sons Ltd., 2014.

[4] D.E. Olivares, A. Mehrizi-Sani, A.H. Etemadi, C.A. Cañizares, R. Iravani, M. Kazerani, A.H. Hajimiragha, O. Gomis-Bellmunt, M. Saeedifard, R. Palma-Behnke, G.A. Jimenez-Estevez, and N.D. Hatziargyriou. Trends in microgrid control. *IEEE Trans. on Smart Grid*, 5(4):1905–1919, 2014.

[5] J. Guerrero, J. C. Vasquez, J. Matas, L. G. de Vicuna, and M. Castilla. Hierarchical control of droop-controlled ac and dc microgrids - a general approach toward standardization. *IEEE Trans. Ind. Electron.*, 58(1):158–172, Jan 2011.

[6] A. Bernstein, L. Reyes-Chamorro, J.-Y. Le Boudec, and M. Paolone. A composable method for real-time control of active distribution networks with explicit power setpoints. Part I: Framework. *Electric Power Systems Research*, 125:254–264, 2015.

[7] Andreas Ulbig, Theodor S. Borsche, and Gran Andersson. Impact of low rotational inertia on power system stability and operation. In *IFAC Proceedings Volumes*, pages 7290 – 7297, 2014.

[8] N. Pogaku, M. Prodanovic, and T. C. Green. Modeling, analysis and testing of autonomous operation of an inverter-based microgrid. *IEEE Trans. Power Electron.*, 22(2):613–625, March 2007.

[9] J. A. Peças Lopes, C. L. Moreira, and A. G. Madureira. Defining control strategies for microgrids islanded operation. *IEEE Trans. Power Syst.*, 21(2):916–924, 2006.

[10] J. Guerrero, M. Chandorkar, T. Lee, and P. Loh. Advanced control architectures for intelligent microgrids - part I: Decentralized and hierarchical control. *IEEE Trans. Ind. Electron.*, 60(4):1254–1262, April 2013.

[11] T. Vandoorn, B. Meersman, J. D. M. De Kooning, and L. Vandevelde. Directly-coupled synchronous generators with converter behavior in islanded microgrids. *IEEE Trans. Power Syst.*, 27(3):1395–1406, Aug 2012.

[12] M. Paolone. Editorial. Sustainable Energy, Grids and Networks, 2015.

[13] L. Reyes-Chamorro. *Real-Time Control Framework for Active Distribution Networks Theoretical Definition and Experimental Validation*. PhD thesis, EPFL, 2016.

[14] K. De Brabandere, B. Bolsens, J. Van den Keybus, A. Woyte, J. Driesen, and R. Belmans. A voltage and frequency droop control method for parallel inverters. *IEEE Trans. on Power Electronics*, 22(4):1107–1115, July 2007.

[15] Q. C. Zhong and Y. Zeng. Universal droop control of inverters with different types of output impedance. *IEEE Access*, 4:702–712, 2016.

[16] J. He and Y. W. Li. Analysis, design, and implementation of virtual impedance for power electronics interfaced distributed generation. *IEEE Trans. on Industry Applications*, 47(6):2525–2538, Nov 2011.

[17] J. Guerrero, J. Matas, L. Garcia de Vicuna, M. Castilla, and J. Miret. Decentralized control for parallel operation of distributed generation inverters using resistive output impedance. *IEEE Trans. Ind. Electron.*, 54(2):994–1004, April 2007.

[18] X. Lu, K. Sun, J. Guerrero, and L. Huang. Soc-based dynamic power sharing method with ac-bus voltage restoration for microgrid applications. In *IECON 2012 - 38th Annual Conference on IEEE Industrial Electronics Society*, pages 5677–5682, Oct 2012.

[19] H. Liang, B. J. Choi, W. Zhuang, X. Shen, A. S. A. Awad, and A. Abdr. Multiagent coordination in microgrids via wireless networks. *IEEE Wireless Communications*, 19(3):14–22, June 2012.

[20] A. Vaccaro, V. Loia, G. Formato, P. Wall, and V. Terzija. A self-organizing architecture for decentralized smart microgrids synchronization, control, and monitoring. *IEEE Trans Ind. Informat.*, 11(1):289–298, Feb 2015.

[21] Christian Rehtanz. *Autonomous Systems and Intelligent Agents in Power System Control and Operation*. Springer, 2003.

[22] Y. S. Foo and H. B. Gooi. Multi-agent system for optimization of microgrids. In *Proceedings of the 8th International Conference on Power Electronics (ECCE Asia)*, May 2011.

[23] C. Zhao, E. Mallada, S. Low, and J. Bialek. A unified framework for frequency control and congestion management. In *Proceedings of the 19th Power Systems Computation Conference (PSCC 2016)*, Jun 2016.

[24] A. G. Expósito, A. Abur, A. de la Villa Jaén, and C. Gómez-Quiles. A multilevel state estimation paradigm for smart grids. *Proceedings of the IEEE*, 99(6):952–976, June 2011.

[25] P. Romano and M. Paolone. Enhanced interpolated-DFT for synchrophasor estimation in FPGAs: Theory, implementation, and validation of a PMU prototype. *IEEE Trans. on Instrumentation and Measurement*, 63(12):2824–2836, December 2014.

[26] A. Derviskadic, P. Romano, M. Pignati, and M. Paolone. Architecture and experimental validation of a low-latency phasor data concentrator. *IEEE Trans. on Smart Grid*, PP(99):1–1, 2016.

[27] S. Sarri, L. Zanni, M. Popovic, J. Y. Le Boudec, and M. Paolone. Performance assessment of linear state estimators using synchrophasor measurements. *IEEE Trans. on Instrumentation and Measurement*, 65(3):535–548, March 2016.

[28] M. Popovic, M. Mohiuddin, D. C. Tomozei, and J. Y. L. Boudec. iPRP - the parallel redundancy protocol for ip networks: Protocol design and operation. *IEEE Trans Ind. Informat.*, 12(5):1842–1854, Oct 2016.

[29] T. Tesfay and J.Y. Le Boudec. Experimental comparison of multicast authentication for wide area monitoring systems. *IEEE Trans. on Smart Grid*, PP(99):1–1, 2017.

[30] K. Christakou, J.-Y. Le Boudec, M. Paolone, and D.-C. Tomozei. Efficient computation of sensitivity coefficients of node voltages and line currents in unbalanced radial electrical distribution networks. *IEEE Trans. on Smart Grid*, 4(2):741–750, June 2013.

[31] L. Reyes-Chamorro, A. Bernstein, J.-Y. Le Boudec, and M. Paolone. A composable method for real-time control of active distribution networks with explicit power setpoints. Part II: Implementation and Validation. *Electric Power Systems Research*, 125:265–280, 2015.

[32] M. Bahramipanah, D. Torregrossa, R. Cherkaoui, and M. Paolone. Enhanced electrical model of lithium-based batteries accounting the charge redistribution effect. In *2014 Power Systems Computation Conference*, pages 1–8, Aug 2014.

[33] F. Sossan, E. Namor, R. Cherkaoui, and M. Paolone. Achieving the dispatchability of distribution feeders through prosumers data driven forecasting and model predictive control of electrochemical storage. *IEEE Trans. on Sustainable Energy*, 7(4):1762–1777, Oct 2016.

[34] L. Reyes-Chamorro, A. Bernstein, M. Paolone, and J.-Y. Le Boudec. A supercapacitor agent for providing real-time power services to the grid. In *Proceedings of the 2015 ISGT LA*, 2015.

[35] E. Scolari, D. Torregrossa, J. Y. Le Boudec, and M. Paolone. Ultra-short-term prediction intervals of photovoltaic ac active power. In *2016 International Conference on Probabilistic Methods Applied to Power Systems (PMAPS)*, pages 1–8, Oct 2016.

[36] G. T. Costanzo, A. Bernstein, L. R. Chamorro, H. W. Bindner, J. Y. Leboudec, and M. Paolone. Electric space heating scheduling for real-time explicit power control in active distribution networks. In *IEEE PES Innovative Smart Grid Technologies, Europe*, pages 1–6, Oct 2014.

[37] S. Papathanassiou, N. Hatziargyriou, and K. Strunz. A benchmark low voltage microgrid network. In *Proceedings of the CIGRÉ Symposium "Power Systems with Dispersed Generation: technologies, impacts on development, operation and performances"*, Apr. 2005, Athens, Greece.

[38] D. E. Olivares, C. A. Cañizares, and M. Kazerani. A centralized energy management system for isolated microgrids. *IEEE Trans. on Smart Grid*, 5(4):1864–1875, July 2014.

[39] F. Garcia-Torres and C. Bordons. Optimal economical schedule of hydrogen-based microgrids with hybrid storage using model predictive control. *IEEE Trans. Ind. Electron.*, 62(8):5195–5207, Aug 2015.

[40] E. Stai, L. Reyes-Chamorro, F. Sossan, J.-Y. Le Boudec, and M. Paolone. Dispatching stochastic heterogeneous resources accounting for grid and battery losses. *IEEE Transactions on Smart Grid*, PP(99):1–1, 2017.

Enrica Scolari (S15) was born in Italy in 1989. She received the M.Sc. (Hons.) degree in Energy Engineering from the University of Bologna, Italy, in 2014. Since 2015, she is pursuing the Ph.D. degree with the Distributed Electrical Systems Laboratory, Swiss Federal Institute of Technology, Lausanne, Switzerland. Her main research interests are solar photovoltaic modelling and forecasting applied to power system.

Andreas M. Kettner (M15) grew up in Sünikon, Switzerland, and attended the Kantonsschule Zürcher Unterland in Bülach, Switzerland. He received the B.Sc. and M.Sc. degrees in Electrical Engineering and Information Technology from the Eidgenössische Technische Hochschule Zürich, Switzerland, in 2012 and 2014, respectively. After working as a development engineer for Supercomputing Systems AG in Zürich, he joined the Distributed Electrical Systems Laboratory at the École Polytechnique Fédérale de Lausanne, Switzerland, where he is pursuing a Ph.D. degree. His research interests include real-time monitoring and control of Active Distribution Networks with particular focus on State Estimation and Voltage Stability Assessment.

Lorenzo Reyes-Chamorro (S13-M16) was born in Santiago, Chile in 1984. He received the B.Sc. degree in electrical engineering from the University of Chile, Santiago, Chile in 2009 and the Ph.D. degree at the Swiss Federal Institute of Technology of Lausanne (EPFL), Lausanne, Switzerland in 2016. He is currently a postdoctoral fellow at the Distributed Electrical System Laboratory, EPFL and an ad-honorem Professor at the Austral University of Chile, Valdivia, Chile.

Andrey Bernstein (M15) received the B.Sc. and M.Sc. degrees in Electrical Engineering from the Technion - Israel Institute of Technology in 2002 and 2007 respectively, both summa cum laude. He received the Ph.D. degree in Electrical Engineering from the Technion in 2013. Between 2010 and 2011, he was a visiting researcher at Columbia University. During 2011-2012, he was a visiting Assistant Professor at the Stony Brook University. From 2013 to 2016, he was a postdoctoral researcher at the Laboratory for Communications and Applications of Ecole Polytechnique Federale de Lausanne (EPFL), Switzerland. Since October 2016 he has been a Senior Scientist at the National Renewable Energy Laboratory, Golden, CO, USA. His research interests are in the decision and control problems in complex environments and related optimization and machine learning methods, with particular application to intelligent power and energy systems. Current research is focused on real-time control of power distribution systems with high penetration of renewables by using online optimization methodology, and machine learning methods for grid data analytics.

Niek J. Bouman received the B.Sc. degree and M.Sc. degree in Electrical Engineering from Universiteit Twente, Enschede, NL, in 2005 and 2007, respectively, and the Ph.D. degree in Mathematics from Universiteit Leiden, Leiden, NL, in 2012. From 2008 to 2012, he was with the cryptology group of Centrum Wiskunde & Informatica (CWI), Amsterdam, NL, where he worked on quantum cryptography and (quantum) information theory. During 2013-2014, he was an independent research consultant for a crypto-currency-analysis project. From 2014 to 2016 he was with the Laboratory for Communications and Applications of EPFL, Lausanne, Switzerland. Currently, he is affiliated to Eindhoven University of Technology, as a researcher in cryptography, multiparty computation and privacy-preserving data mining. His research interests include information and communication theory, cryptography and multiparty computation, statistical inference, as well as applied research related to grid control and the Internet of Things.

Benoit Cathiard was an intern at the Laboratory for Communications and Applications of the École Polytechnique Fédérale de Lausanne during 2016.
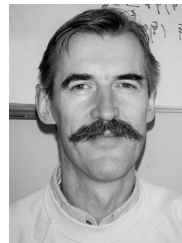
Jean-Yves Le Boudec (F04) is professor at EPFL and fellow of the IEEE. He graduated from Ecole Normale Supérieure de Saint-Cloud, Paris, where he obtained the Agrégation in Mathematics in 1980 and received his doctorate in 1984 from the University of Rennes, France. From 1984 to 1987 he was with INSA/IRISA, Rennes. In 1987 he joined Bell Northern Research, Ottawa, Canada, as a member of scientific staff in the Network and Product Traffic Design Department. In 1988, he joined the IBM Zurich Research Laboratory where he was manager of the Customer Premises Network Department. In 1994 he became associate professor at EPFL. His interests are in the performance and architecture of communication systems and smart grids. He co-authored a book on network calculus, which forms a foundation to many traffic control concepts in the internet, an introductory textbook on Information Sciences, and is also the author of the book "Performance Evaluation".

Mario Paolone (M07-SM10) received the M.Sc. (Hons.) and Ph.D. degrees in electrical engineering from the University of Bologna, Bologna, Italy, in 1998 and 2002, respectively. In 2005, he was nominated Assistant Professor in power systems at the University of Bologna, where he was with the power systems laboratory until 2011. He is currently an Associate Professor at the Swiss Federal Institute of Technology, Lausanne, Switzerland, chair of the Distributed Electrical Systems Laboratory. He was the Co-Chairperson of the Technical Committee of the 9th edition of the International Conference of Power Systems Transients (2009) and Vice-Chair and Chair of the 19th and 20th Power Systems Computation Conference in 2016 and 2018 respectively. He is author or co-author of more than 220 scientific papers published in reviewed journals and international conferences. He is the Editor-in-Chief of the Elsevier journal Sustainable Energy, Grids and Networks and the Head of the Swiss Competence Center for Energy Research "FURIES". His research interests include power systems with particular reference to realtime monitoring and operation of active distribution networks, integration of distributed energy storage systems, power system protections and power system transients. In 2013, he received the IEEE EMC Society Technical Achievement Award.