A composable method for real-time control of active distribution networks with explicit power setpoints. Part II: Implementation and validation

Lorenzo Reyes-Chamorro [b,*,1], Andrey Bernstein [a,1], Jean-Yves Le Boudec [a], Mario Paolone [b]

[a] *Laboratory for Communications and Applications 2, EPFL, CH-1015 Lausanne, Switzerland*
[b] *Distributed Electrical Systems Laboratory, EPFL, CH-1015 Lausanne, Switzerland*

ARTICLE INFO

ABSTRACT

In this second part, we evaluate the performances of our control framework by applying it to a case study that contains a minimum set of elements allowing to show its applicability and potentials. We show how the computation of the *PQt* profiles, belief functions, and virtual costs can be synthesized for generic resources (i.e., dispatchable and stochastic generation systems, storage units, loads). The metrics of interest are: quality-of-service of the network represented by voltages magnitudes and lines current magnitudes in comparison with their operational boundaries; state-of-charge of electric and thermal storage devices; proportion of curtailed renewables; and propensity of microgrid collapse in the case of renewables overproduction. We compare our method to two classic ones relying on droop control: the first one with only primary control on both frequency and voltage and the second one with an additional secondary frequency control operated by the slack device. We find that our method is able to indirectly control the reserve of the storage systems connected to the microgrid, thus maximizing the autonomy in the islanded operation and, at the same time, reducing renewables curtailment. Moreover, the proposed control framework keeps the system in feasible operation conditions, better explores the various degrees of freedom of the whole system and connected devices, and prevents its collapse in case of extreme operation of stochastic resources. All of these properties are obtained with a simple and generic control framework that supports aggregation and composability.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

In this second part of the paper, we discuss the implementation aspects and evaluate the performance of the control framework, which we henceforth refer to as *Commelec* (which stands for the joint-operation of Communication and Electricity systems). This assessment is done by using a suitably developed simulation environment. We consider a case study that makes reference to the low voltage microgrid benchmark defined by the CIGRÉ Task Force C6.04.02 [1], connected to a generic medium voltage feeder that contains the minimum number of elements that allow us to show the applicability and potentials of the proposed control framework. In particular, while the formal description of the framework for controlling the grid using power setpoints is presented in Part I, here we show how to specifically implement the request/advertise messages between agents, how we can derive the *PQt* profiles, belief and cost functions of the resources, and how the grid agent computes the resources setpoints and aggregates their internal elements.

The considered case study exhibits the following characteristics: (i) the system is in islanded condition, (ii) the slack bus is provided by the larger storage system (ESS), (iii) storage is distributed in both low and medium voltage, (vi) thermal loads (water boilers) are used as virtual storage, and that (v) the randomness comes from the loads absorption patterns and solar irradiation. For the latter, we used a high time-resolution profile (sampled each 50 ms) obtained from the measurements on solar panels in the authors' laboratory.

A challenge in such a system is that most of the inertia comes from storage and thermal loads rather than rotating machines; it

* Corresponding author at: EPFL STI IEL DESL, ELL037, Station 11, CH-1015 Lausanne, Switzerland. Tel.: +41 21 69 37369.
E-mail addresses: lorenzo.reyes@epfl.ch (L. Reyes-Chamorro), andrey.bernstein@epfl.ch (A. Bernstein), jean-yves.leboudec@epfl.ch (J.-Y. Le Boudec), mario.paolone@epfl.ch (M. Paolone).
[1] Both authors contributed equally in this research work.

Table 1
Electrical parameters of the lines composing the MV and LV lines

Type	MV1	LV1	LV2	LV3	LV4	LV5	LV6
Resistance [Ω/km]	3.938	0.284	0.497	3.690	1.380	0.871	0.822
Reactance [Ω/km]	1.969	0.083	0.086	0.094	0.082	0.081	0.077
Susceptance [μs/km]	2.780	0	0	0	0	0	0
Ampacity [A]	25	170	120	31	60	73	140

is precisely the goal of our real-time control method, to overcome this difficulty in the presence of extremely volatile resources (e.g., PVs).

In order to assess its performance, we used the following metrics: the distances of node voltages and line currents to their operational limits, the state-of-charge of electric and thermal storage devices, the proportion of curtailed renewables, and the robustness against system collapse in case of overproduction from renewables.

We compare our method to two classic ones that rely on droop control: the first only with *primary control* on both frequency and voltage and the second with an additional *secondary frequency control* at the slack device (see, e.g., [2]). We find that our method is able to indirectly control the reserve of the storage systems, thus maximizing the autonomy of the islanded operation. It reduces the curtailment of renewables, compared to the droop based methods, and it is able to implicitly identify local power compensation. Further, it keeps the system in feasible operation conditions and better explores the various degrees of freedom of both network and energy resources. Most importantly, it prevents system collapse in case of overproduction of renewables, in contrast to the droop control strategies.

Further, we show that the properties of Commelec are fundamental in the case of *inertia-less grids* associated with the penetration of energy conversion systems that do not have any rotating mass (e.g., PV plants) or other conversion systems interfaced to the grid with power converters. Indeed, in the cases where these energy conversion systems represent the majority of the electricity supplying means, the control strategies have to be re-thought (e.g., [3]). In this respect, as the proposed method does not rely on any shared signals (i.e., frequency), it can inherently account for the control of inertia-less grids.

All of these characteristics are obtained in real-time with a simple and generic framework; the specific properties of electric and thermal resources are known only by their local agents, whereas grid agents are generic and independent of the specific resources they control. As introduced in Part I, a key property is composability: an entire grid can be viewed as a single generic resource, the details of which need not be known by the higher-level grid agent. In this part of the paper, we also evaluate the effect of the simplifications resulting from the aggregation process, and we find that it is essentially negligible.

The structure of this second part is the following. In Section 2, we present the case to be studied, the simulation environment, the related control algorithms, the profiles' data, and the performance metrics. In Section 3, we define the different resource agents and how they manage their exchanged messages. In Section 4, we present the simulation results. A discussion section and a conclusion follow.

2. Case study

In this section, we present a case study where the proposed control framework is implemented. To show the applicability of the proposed framework, we have selected a closed system that contains all types of agents described in Part I. In order to evaluate its

performance, we implemented a generic event-driven simulation environment in Matlab®.

2.1. System details

We consider a 0.4 kV LV network that includes (i) distributed generation composed of photovoltaic plants (PVi) and a hydraulic microturbine (μH), (ii) a storage system represented by a battery (ESS1), (iii) uncontrollable loads (ULi) and (iv) controllable loads (WBi) modelled as water boilers all capable of deploying explicit control setpoints. The topology and parameters of this LV grid are taken from [1]. As typically used in a microgrid (MG) setup, we assume that all the generation/storage units connected to the LV MG are interfaced with the grid through power electronic devices [4].

To show the interaction between different grids, the MG is connected to a 20 kV MV distribution system that interconnects (i) a large battery storage system (ESS), (ii) a combined heat and power generator interfaced with the MV grid by means of a synchronous generator (SG) and (iii) an industrial uncontrollable load (UL).

The corresponding electrical diagram for the case study is presented in Fig. 1(a).

To illustrate the mapping between physical subsystems and agents, we consider the hierarchical agents setting shown in Fig. 1(b) where the microgrid agent (LVGA) is in charge of the resources in the LV network, whereas the medium voltage grid agent (MVGA) is in charge of the ones in the MV network and the LVGA. In the terminology of Part I of the paper, the LVGA is an internal GA, while the MVGA is a root GA.

The line parameters used for the network are presented in Table 1.

We use the base system and the voltage bounds presented in Table 2(a), while the parameters of the MV/LV transformer used in our case study are shown in Table 2(b). We use a conventional transformer model as in [5].

2.2. Control methods

We perform a comparison between the following control methods.

(i) The Commelec architecture of Fig. 2(a). We show in the following sections how we implement our framework in this case. In addition, in order to validate the composability property we performed a simulation of the "flat" setting of agents shown in Fig. 2(b).

(ii) The droop control method, with only a *primary control* at each device capable of modifying power setpoints (i.e., all with the exception of ULs). In the slack resource, the output frequency is calculated using the conventional droop control strategy, assuming a null inertia (as it is the case of ESS). This is the

Table 2
System and transformer parameters

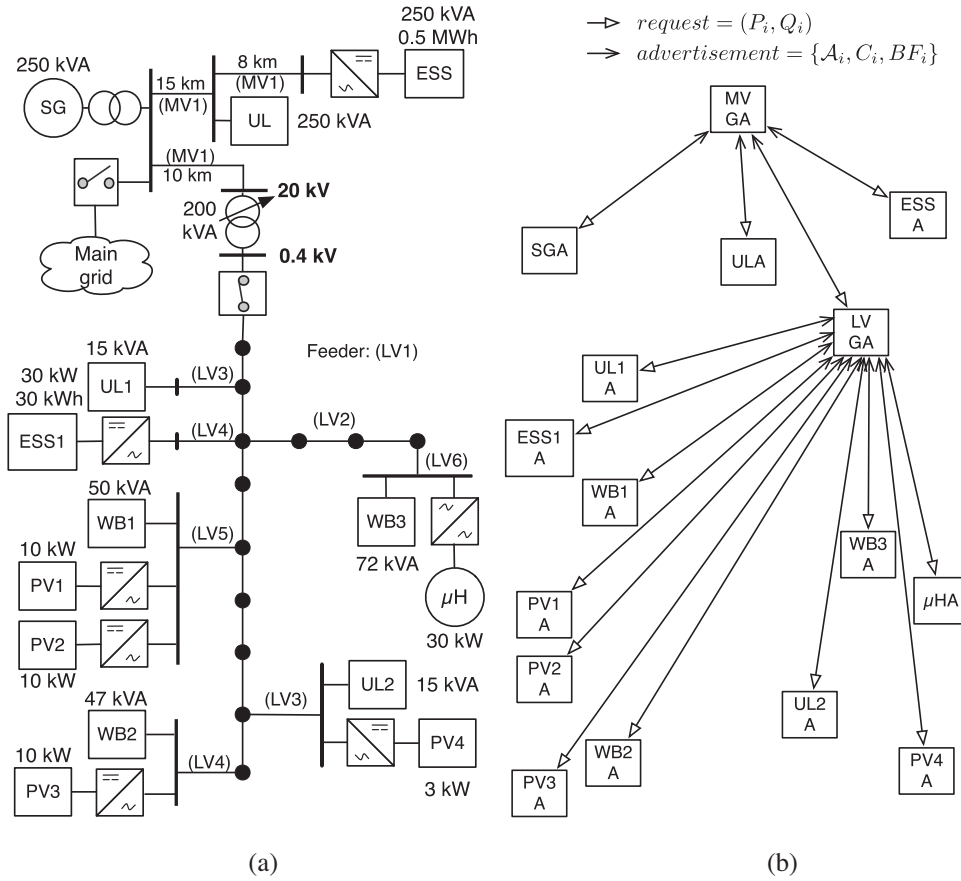| MV base	LV base	Base power	V bounds	S_r	V_k	r_k
20 kV	0.4 kV	1 MVA	0.9–1.1 pu	400 kVA	4%	1%

Fig. 1. The electrical network and agents for the case study. (a) Microgrid. (b) Agents.

signal that will be used for all the other resources to compute their power production. As a result, the frequency is given by $f = f_0 - m_f(P - P_0)$, where f_0 is the rated frequency (here 50 Hz), m_f is the slope, and P_0 is a parameter. The corresponding frequency and voltage droop curves in other resources were set to

$$P = \frac{P_0 - (f - f_0)}{m_f},$$
$$Q = \frac{Q_0 - (V - V_0)}{m_V},$$

where V is the measured voltage magnitude and f the measured frequency, $V_0 = 1$ pu is the rated voltage, m_V is the curve slope, and Q_0 is the reactive power when $V = V_0$. It is worth noting that the droop parameters are in general different for each resource. We have selected the droop parameters for the resources using typical values adopted in the literature (e.g. [2]). The selected droop parameters are shown in Table 3.

(iii) The droop control method as above, with additional *secondary frequency control* at the slack device. In particular,

the frequency droop curve in the slack resource was set to

$$f(t) = f_0 - m_f(P - P_0) + \left(\frac{1}{T_i}\right) \int_{t_0}^{t^-} (f_0 - f(\tau)) d\tau,$$

where $T_i = 50$ sec is the chosen integration constant.

2.3. Profiles data

We chose a simulation scenario for all three control methods (including initial conditions), where we could simulate the case of overall overproduction in the grid from renewables (essentially PV) with minimum load consumption. This scenario is adopted to challenge both control methods to deal with a system characterized by a low margin of controllability. For this purpose, the scenario has the following initial conditions:

- All batteries are close to their maximum stored energy capacity. In particular, the initial state of charge (*SoC*) of both MV and LV battery was set to 0.9.
- The boilers are undercharged, with initial state of 2.5 kWh.
- There is a high production from PVs, at a partially sunny day, thus representing high irradiation variability.
- The loads in the LV grid have zero-consumption profiles, whereas the MV load uses a dynamic profile representing changes with time resolution of 1 min.
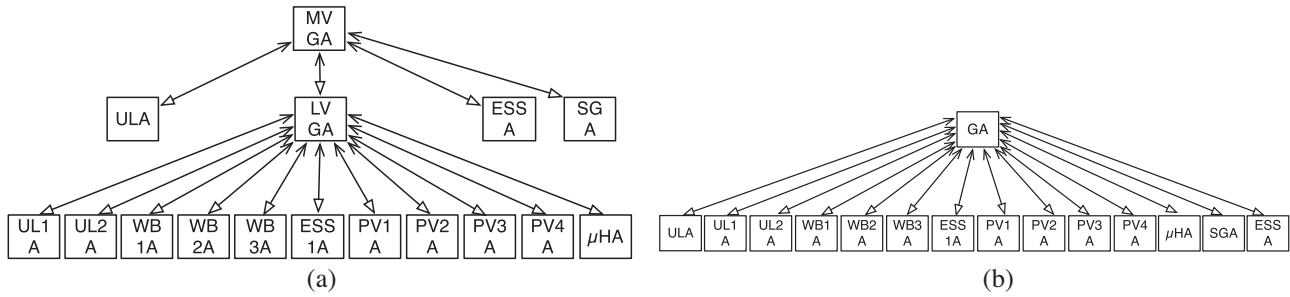
Table 3
Droop parameters.

Resource	f_0 [Hz]	P_0 [pu]	m_f [Hz/pu]	V_0 [pu]	Q_0 [pu]	m_V [pu]
ESS/ESS1	50	0	−0.5	1	0	−0.04
PVi/WBi	50	0.5	−1	1	0.5	−0.08
SG/μH	50	0.5	−0.8	1	0.5	−0.08

Fig. 2. Agents architecture for the case study. (a) Hierarchical, (b) flat.
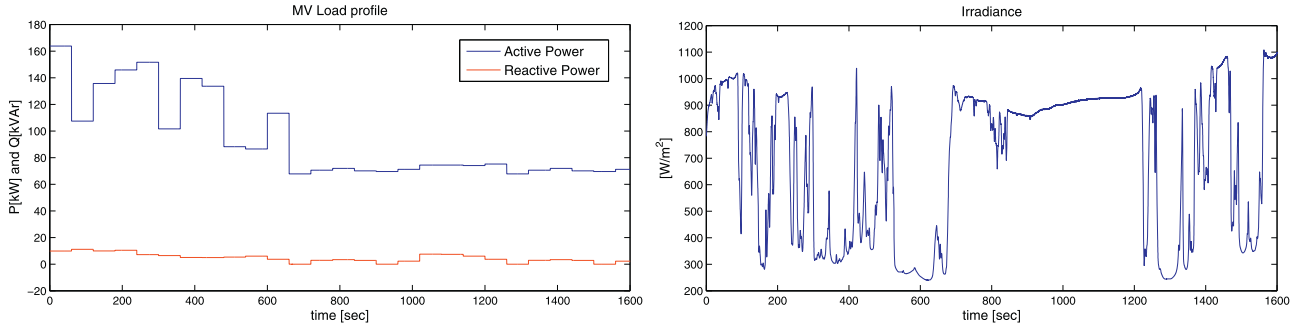

Fig. 3. Sources of uncertainty in the case study: UL load profile and solar irradiance.

In our case study, there are two sources of uncertainty: the MV-load (UL) power consumption and the solar irradiance (shown in Fig. 3(a) and (b), respectively). We assume that all the PV plants are exposed to the same irradiance profile. The load consumption is characterized by a dynamic behaviour and a low value from minute 11 onwards, whereas the solar irradiance data is characterized by a highly volatile profile due to the passage of clouds. The irradiance data corresponds to real measurements performed in the authors' laboratory ($46°31'06.20''$ N, $6°33'54.56''$ E) on November the 15th, 2013. The sampling period used to take the data was 50 ms. The used profile is shown in Fig. 3(b). The above quantities represent the forcing functions of the targeted case study. Lastly, we use the weights shown in Table 4 for Commelec simulations.

2.4. Performance metrics

In order to assess the performance of the control methods, we use the following metrics: (i) the distances of node voltages and line currents to their limits, representing the quality of supply and the operational margins of the system; (ii) the state of charge of electric and thermal storage devices, representing the reserve of the system; (iii) the proportion of curtailed renewables; and (iv) the robustness of the method against system collapse.

3. Resources models and agents

As anticipated, we consider the following resources: (i) *energy storage device* (specifically, a *battery*), (ii) *synchronous generator*, (iii) *PV generator* and (iv) *controllable and uncontrollable loads*.

Table 4
Simulation parameters.

α_{LV}	α_{MV}	$\omega_{PV_i}, \omega_{UL_i}, \omega_{\mu H}, \omega_{UL}$	$\omega_{WB_{1-3}}$	ω_{WB_2}	ω_{ESS_1}	ω_{ESS}	ω_{SG}	$\omega_{(P_0,Q_0)}$
0.1	$5E-5$	1	$1E-6$	$1E-8$	$1E-3$	$1E-5$	0.01	100

Depending on their nature and/or internal characteristics, these resources have various degrees of controllability, from *fully controllable* resources (e.g., the battery) to *non-controllable* resources (e.g., uncontrollable load). The controllability of the resource has a considerable effect on the design of the corresponding resource agent (RA).

RAs are pieces of software usually deployed on a computer, a processor or a microcontroller installed in the vicinity of the resource. For instance, in a generation or storage unit, the RA can be implemented within its controller, whereas a load agent can be installed in a building computer to monitor and/or control its aggregated power consumption. RAs might have a simplified or sophisticated view of the internal behaviour of their resources as a function of the RA developer. The better is the resource model, the more accurate advertisement messages will be sent to the GA, and the better would be its decision.

Recall that RA communication messages refer to the power flows at the point of connection with the grid, thus converters are always considered as part of a resource. As a converter can be used by all kind of resources, we first present a general approach for its model as an interface with the grid. Further, we present how to implement RAs in detail, specifically, how they manage the requests and produce the advertisements.

3.1. Power converter

First, we consider that the admissible area of operation of power converters can be modelled with three general constraints:

(i) *The PQ capability curve of the converter*, which we consider in this paper to be given by $\sqrt{P^2 + Q^2} \leq S_r$, with S_r the rated power of the converter and (P,Q) the AC powers of the converter. Alternatively, it could be considered that a converter is constrained by its nominal current I_n. In such a case the capability curve is given by $\sqrt{P^2 + Q^2} \leq VI_n$, where V is the voltage magnitude at the connection point. As V can be externally affected, the resulting power constraint has an uncertainty

that has to be reflected in the belief function. For the sake of simplicity, we work in this paper with power constraints.

(ii) *The power factor constraint*, given by

$$\left| \frac{P}{\sqrt{P^2 + Q^2}} \right| \geq \cos_{\min}(\phi), \tag{1}$$

where $\cos(\phi)$ stands for the power factor of the converter and ϕ the phase-shift between voltage and current phasors. This constraint is relevant, for instance, in the case of PV converters that are required to operate with a minimum power.

(iii) *Unidirectional or bidirectional converter*, depending on the nature of the resource. For instance, a grid-tie PV usually cannot absorb active power, thus $P \geq 0$. In the case of an energy storage system, we consider a unique bidirectional device for charge and discharge.

For simplicity, we assume a constant efficiency (η) to account for the effect on the DC power (P_{dc}), depending on the power flow direction:

$$P = \begin{cases} \eta P_{dc}, & \text{if } P_{dc} \geq 0, \\ P_{dc}/\eta, & \text{if } P_{dc} < 0. \end{cases} \tag{2}$$

3.2. Energy storage system

For concreteness, we consider the case in which the ESS is represented by a battery. (However, the concepts and methods can be easily extended to any kind of ESS.)

Implementation of setpoints. In order to implement a requested power setpoint, the ESS agent (ESSA) needs a model to compute the internal limits this resource must respect for the next time step. In this paper, we consider that such an agent uses a simple model that can sufficiently represent the dynamic behaviour of the storage system in the considered time frame.[2] In particular, assuming that the state of charge (*SoC*) is fixed between two setpoint implementations, we can express the model of the battery as a simple time-varying resistance R_t that is a function of the *dc* current and voltage measurements of the battery array. This approximation is reasonable if frequent battery setpoint variations are deployed, enabling a pseudo-continuous computation of R_t.

On the contrary, in our simulation environment we use a two time constants (TTC) cell model (e.g., [7]) to simulate the internal behaviour of the battery.

Upon receiving a new setpoint request at time t, the ESSA computes

$$R_t = \frac{\Delta V^{dc}}{\Delta I^{dc}} = \frac{V_t^{dc} - V_{t-\Delta t}^{dc}}{I_t^{dc} - I_{t-\Delta t}^{dc}} \tag{3}$$

where ΔV^{dc} and ΔI^{dc} are the step changes in *dc* voltage and current measured in the resource at two consecutive requests instants. (Note that if $\Delta I^{dc} = 0$, R_t will not change). Consequently the ESSA can compute the internal electromotive force of the bank as $E_t = R_t I_t^{dc} + V_t^{dc}$. Then, by means of this extremely simple model, and considering the limitations on V^{dc} and I^{dc} given by the storage specifications (V_{min}, V_{max}, I_{min} and I_{max}), ESSA computes the *dc* power bounds for the resource as

$$P_{\min}^{dc} = \max \left(\frac{V_{\max}(E_t - V_{\max})}{R_t}, (E_t - R_t I_{\min})I_{\min} \right), P_{\max}^{dc}$$

$$= \min(P_{\max}^{V_{dc}}, P_{\max}^{I_{dc}})$$

$$P_{\max}^{V_{dc}} = \begin{cases} E_t^2/4R_t, & \text{if } \frac{E_t}{2} > V_{\min}, \\ \dfrac{V_{\min}(E_t - V_{\min})}{R_t}, & \text{otherwise,} \end{cases}, P_{\max}^{I_{dc}}$$

$$= \begin{cases} E_t^2/4R_t, & \text{if } \frac{E_t}{2R_t} > I_{\max}, \\ (E_t - R_t I_{\max})I_{\max}, & \text{otherwise.} \end{cases}$$

The above *dc* power bounds[3] are combined with the converter model in (2) to compute the *ac* active power bounds. Finally, it projects the requested setpoint onto the set of constraints defined by these bounds and the converter constraints.

PQt **Profile**. As the constant *SoC* assumption is still valid, the power bounds for this resource advertised within the *PQt* profile are fully specified by the aforementioned process. In this respect, a *PQt* profile slice for a given time step, $\mathcal{A}_b$, is shown in Fig. 4(a).

Belief function. As storage devices are highly controllable, we assume an ideal belief, namely, $BF_b(u_b) = \{u_b\}$ for any u_b belonging to the *PQt* Profile.

Virtual cost. The role of the virtual cost function is to measure the tendency of the storage agent to stay within particular zones of the *PQt* profile. In this paper, we consider that ESSA tends to steer the *SoC* to a certain target value that represents a suitable (admissible) internal state of the storage obtained from a long-term scheduler.[4] Therefore, if the current *SoC* is larger than a target value, the ESS prefers to be discharged, so the agent advertises a negative cost for discharging (positive P) and a positive cost for charging (negative P). This situation is reversed for the case when the *SoC* is lower than the target value. If the *SoC* is equal to the target, the cost will become zero, as the agent shows no preference of the ESS to be charged or discharged. We assume that the current *SoC* is measured by the resource using the *SoC* computation presented in [8].

As an example, the following polynomial function can be used

$$C_b(P, Q) = k\Delta SoC \cdot \left(a_b P^2 + \frac{b_b}{\Delta SoC}P + c_b \right) \cdot P, \tag{4}$$

where $\Delta SoC = SoC_t - SoC$, SoC_t is the target *SoC*, and a_b, b_b, c_b and k are positive constants. This function is chosen so that it presents (i) a positive *cost* when going in the opposite direction of the target *SoC*, and a negative cost (incentive) when heading towards the target; and (ii) a higher *price* (dC_p/dP) for a higher power at constant *SoC*. An illustration of this function is shown in Fig. 4(b) for different values of *SoC*. For example, when $\Delta SoC > 0$, the cost for charging is positive with a steep slope and the cost for discharging is negative with gentle slope varying with asked power. It should be noted that the cost for reactive power for energy storage systems is considered to be zero.

3.3. Photovoltaic plant

Implementation of setpoints. Using measurements in the resource, we assume that the PV agent (PVA) can obtain the current maximum admissible power production $P_{pv}^{\max}$. Then, the PVA controls its resource to set the request u_{pv} projected to the admissible set defined by this bound and the converter limits from Section 3.1. Afterwards, as described in Part I, the resource tries to deploy such a setpoint. The variation between the requested setpoint u_{pv} and the actual setpoint x_{pv} is represented by the belief function of the PVA.

---

[2] However, more complex models (e.g., [6]) can be easily made compatible with the proposed framework.

[3] Note that the arguments of $P_{\min}^{dc}$ are always negative.

[4] For electrochemical storage systems, this scheduler may take into account their state-of-health and consequent life.
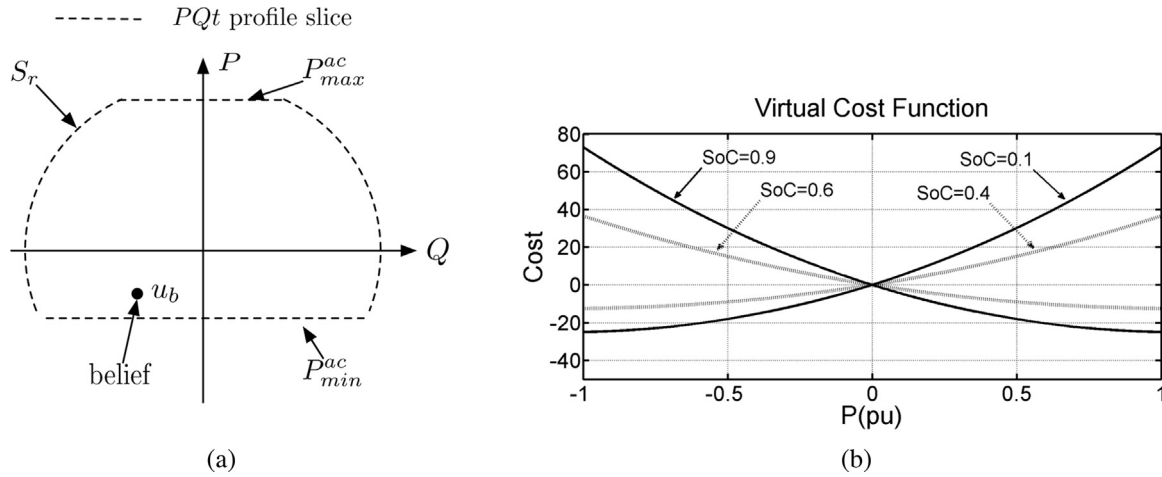
Fig. 4. PQt profile, belief function and virtual cost for ESS agent.

PQt **Profile**. By means of a forecasting tool, and using the converter model (2), the PVA computes the maximum power production at time $t \in [t_0, t_0 + T]$, $P_{pv}^f(t)$, that can be maintained for any $t' \in [t, t + T]$. As typically for grid-tie PV converters, we assume a constraint on the reactive power production, given by a minimum power factor (1). For time t, a slice of the PQt profile shown in dashed lines in Fig. 5(a).

Belief function. In order to advertise the uncertainty of the solar resource, we assume that the active power production might decrease from the requested setpoint, u_{pv}, with a predicted maximum variation ΔP_{pv}^{max}. The value of ΔP_{pv}^{max} is determined from the worst-case error of the employed forecasting tool. As the reactive power is controlled by the converter, the belief of Q production is restricted only by its relation with P and the constraints that define the PQt profile. Hence, $BF_{pv}(u_{pv})$ can be represented as a line that starts in $u_{pv} = (P, Q)$ and finishes in $u'_{pv} = (P', Q')$, with $P' = P - \Delta P_{pv}^{max}$ and

$$
Q' = \begin{cases} \max \left\{ -P' \dfrac{\sqrt{1 - \cos_{min}^2(\phi)}}{\cos_{min}(\phi)}, Q \right\}, & \text{if } Q < 0, \\ \min \left\{ P' \dfrac{\sqrt{1 - \cos_{min}^2(\phi)}}{\cos_{min}(\phi)}, Q \right\}, & \text{otherwise.} \end{cases}
$$

An example of $BF_{pv}(u_{pv})$ is shown in Fig. 5(a).

Virtual cost. We assume that the PVA seeks to maximize the active power production and minimize the reactive power. Therefore, an example of the advertised virtual cost function is given by $C_{pv}(P, Q) = -a_{pv}P + b_{pv}Q^2$, with $a_{pv}, b_{pv} > 0$.

3.4. Synchronous generator

For simplicity, we consider cylindrical rotor machines in both synchronous generators (SG and μH), whereas the agent uses the basic model for generator (both equivalent circuit and relevant capability curves) as in [5], and we assume that they are interfaced to the network through an appropriate transformer (in the case of SG) or power converter (in the case of μH). We present this section by making specific reference to SG, but is also applicable to μH. Furthermore, we assume that the inertia of such machines is small enough to express its behaviour by using algebraic equations. This assumption appears reasonable in view of the typical capacity of synchronous machines connected to LV grids and as we assume that the slack resource ESS is connected through a power converter with which frequency can be kept constant. As for the MV grid, it is a reasonable assumption that cannot be extended in general. Indeed, as the proposed Commelec control method does not require the use of the frequency signal, the control of the slack resource can
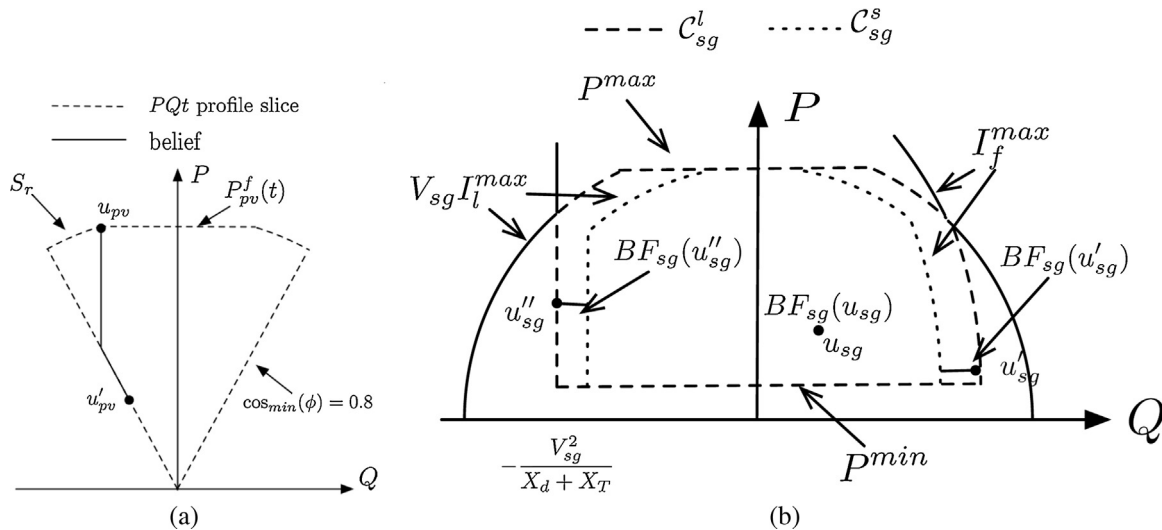

Fig. 5. PQt profile and belief function for PV and SG agents. (a) PVA. (b) SGA.

be astatic. For this reason, the above simplification only improves droop simulation results.

Implementation of setpoints. When receiving a request, the SG agent (SGA) computes the current limits of the resource by using the measurement of the voltage in the connection point V_{sg}. These limits correspond to the well-known capability curves of a synchronous machine defined by maximum and minimum active power P^{max} and P^{min}, maximum SG field-current I_f^{max}, maximum line current I_l^{max}, and stability limit. Further, the SGA commands the implementation of the projection to the capability curves.

PQt Profile. As the bounds of this resource are dependent on V_{sg}, which is in turn dependent on external variables, the prediction of the limits in the next time step is a complex task. Instead, the SGA advertises the *largest possible* set of power setpoints, $\mathcal{C}_{sg}^l$, taking into account all feasible values of V_{sg}.

Belief function. As mentioned above, due to changes in V_{sg}, the boundaries of the capability curves may vary at a given time slot. Thus, some setpoints, in this case the nearest to the bounds, might be shifted to the *smallest possible* set of setpoints, $\mathcal{C}_{sg}^s$, taking into account all feasible values of V_{sg}. Thus, the belief of a given request u_{sg} is

$$BF_{sg}(u_{sg}) = \begin{cases} u_{sg}, & \text{if } u_{sg} \in \mathcal{C}_{sg}^s, \\ \bigcup_{C_{sg}(V_{sg})} \mathcal{P}_{\mathcal{C}_{sg}(V_{sg})}\{u_{sg}\}, & \text{if } u_{sg} \in \mathcal{C}_{sg}^l \backslash \mathcal{C}_{sg}^s, \end{cases}$$

where $\mathcal{P}_{\mathcal{C}_{sg}(V_{sg})}\{\cdot\}$ with $V_{sg} \in [V_{min}, V_{max}]$ is the Euclidean projection to $\mathcal{C}_{sg}(V_{sg})$. An illustration of both *PQt* profile and belief is presented in Fig. 5(b).

Virtual cost. To express the virtual cost, we consider that the SGA operates the resource in order to maximize its overall efficiency. As the efficiency of the turbine given an electrical produced power, $\eta(P)$, plays the most important role in the overall efficiency, we define the virtual cost as $C_{sg}(P, Q) = a_{sg}(1 - \eta(P))$. As an example, the cost function for μHA is given in Fig. 6, with a_{sg} a positive constant.

3.5. Water boiler

We consider thermal controllable loads such as *water boilers* (WB). Each controllable load is considered as a *single* boiler capable of estimating its own thermal state. The approach can be extended to *distributed* controllable loads (e.g., [9]), but this is out of the scope of this paper. Also, a more sophisticated definition of a controllable load agent representing space heating has been presented in [10].

In this section, we consider only active power, with $Q \equiv 0$ throughout.

Implementation of setpoints. We assume that the internal controller of the WB is capable of any active power in $[0, P_{wb}^{max}]$. From our sign convention it follows that $P = 0$ represents the case when the heating device is off, and $P = -P_{wb}^{max}$ represents the case when the heating device is working at full power. Next, we assume that the thermal state of the boiler is represented by the total energy stored in it at time t, given by

$$\mathcal{E}(t) = \int_{\tau=0}^{t} (P_{in}(\tau) - P_{out}(\tau)) d\tau,$$

where $P_{in}(t)$ is the absolute value of the power injected into the system, and $P_{out}(t)$ is the absolute value of the power drawn from the system. The process $\{P_{out}(t)\}$ is the source of uncertainty in this resource, as it is affected by nature and demand patterns of the users of the boiler. The process $\{P_{in}(t)\}$, on the contrary, is controlled by the WB agent (WBA).

The constraints on the energy are given by four parameters $\mathcal{E}^{min} < \mathcal{E}_{margin}^{min} < \mathcal{E}_{margin}^{max} < \mathcal{E}^{max}$.[5] Given a requested setpoint $P \in [-P_{wb}^{max}, 0]$, WBA commands the internal controller to maintain $P_{in}(t) = -P$ as close as possible. Whenever $\mathcal{E}(t) < \mathcal{E}^{min}$, it switches the setpoint to the maximal heating power (namely, to $P_{in}(t) = P_{wb}^{max}$), until $\mathcal{E}(t) \geq \mathcal{E}_{margin}^{min}$. Then, it switches back to the original request, until the energy constraint is violated again. A similar process is assumed when $\mathcal{E}(t) > \mathcal{E}^{max}$. Fig. 7(a) shows this concept.

PQt Profile. We assume that the stored energy is constant until the next request implementation (which is satisfied in view of the large difference between the period of computation of the setpoints T and the resource's time constants). Hence, the *PQt* profile is specified by the interval $[0, -P_{wb}^{max}]$. An example of *PQt* profile is shown in Fig. 7(b).

Belief function. In contrast to a regular storage device, the WB load can be highly uncertain. To account for this uncertainty, we assume that the WBA has a forecasting tool to predict the load profile. Let $[\underline{P}_{out}^f(t), \bar{P}_{out}^f(t)]$ denote the confidence interval of the forecast at time t. To compute the belief set at time t for a given setpoint P, the WBA first computes the worst cases $(\hat{\mathcal{E}}^{min}(t'), \hat{\mathcal{E}}^{max}(t'))$ of the estimated energy at times $t' \in [t, t + T]$ by using the confidence interval. The belief $BF_{wb}(P)$ is then given by either $\{P, 0\}$, $\{-P_{wb}^{max}, P\}$, or $\{-P_{wb}^{max}, P, 0\}$ depending on whether, for some t', $\hat{\mathcal{E}}^{max}(t') > \mathcal{E}^{max}$ or $\hat{\mathcal{E}}^{min}(t') < \mathcal{E}^{min}$, or both. Otherwise, if no violation occurs, the belief is $BF_{wb}(P) = \{P\}$. An example of belief function is given in Fig. 7(b).

Virtual cost. Similarly to the ESS agents, we assume that the basic goal of WBA is to keep the stored energy at a certain target level $\mathcal{E}_{target}$. Therefore, the virtual cost function advertised by WBA is similar to that advertised by ESSA as shown in Fig. 4(b), but centred around the forecasted value of the demand given by $P_{center} = -P_{out}^f(t)$.

3.6. Uncontrollable load

Implementation of setpoints. The UL agent (ULA) does not take into account the requested setpoint as it does not have any way to set it.

PQt Profile. We implement the simplest case, where the *PQt* profile is given by $\{x_l^f(t) = (P^f, Q^f)\}_{t=t_0}^{t_0+T}$. Specifically, for each time step, the *PQt* profile is defined by a *single point* $x_l^f(t)$ given by a *demand forecasting tool*.

Belief function. In this paper, we assume that the UL can change to *any* admissible value at any moment. Hence, the belief is considered as the complete area of operation of the UL. We assume that the consumption of the UL is always inside the semi-circle defined by its maximum apparent power S_r (or ρ_{max} in polar coordinates); that is, it can consume active power and to inject or absorb reactive power. With this representation, the belief is defined by $BF_l = \{(\rho, \theta) : \rho \in [0, \rho_{max}], \theta \in [180°, 360°]\}$, as can be seen in Fig. 8.

Virtual cost. As the ULA cannot control its resource, we set the advertized virtual cost to $C_{ul}(P, Q) = 0$. We note that, in our implementation of the grid agent, only the *gradient of the cost* is used by the employed gradient descent algorithm (see Subsection 5.2.1 in Part I for details). Hence, the uncontrollable load agents can have any *constant* cost without influencing the setpoints computation procedure.

For the different resources here presented, we use the parameters shown in Table 5.

---

[5] It is assumed that there are two levels of stored energy margins: $\mathcal{E}_{margin}^{min}$ and $\mathcal{E}_{margin}^{max}$ have been intended to assure an acceptable margin of operation.
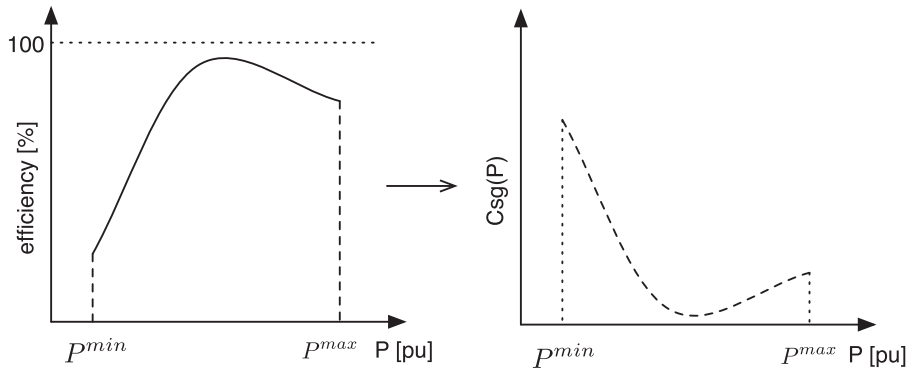
Fig. 6. Efficiency curve of a turbine (left) and cost function of SG agent (right).
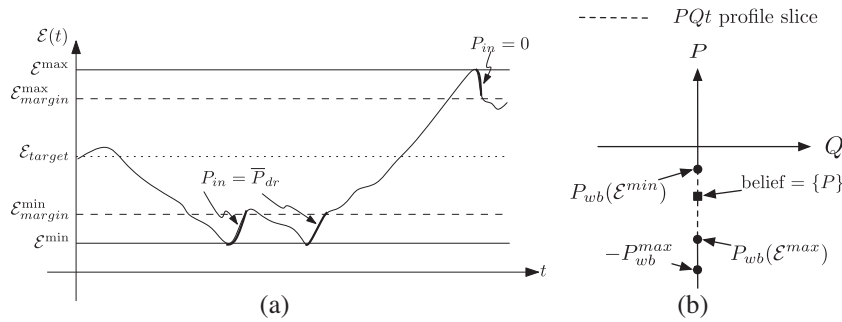

Fig. 7. (a) A possible trajectory of boiler energy as a function of time t. (b) PQt profile and belief function for a WB agent.

Table 5
Resources parameters.

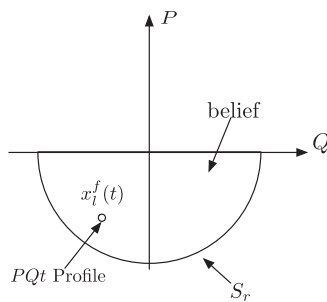| Water boilers				Storage		Synchronous generators					
$\mathcal{E}^{min}$	$\mathcal{E}^{max}$	$\mathcal{E}^{min}_{margin}$	$\mathcal{E}^{max}_{margin}$	η_{ESS}	η_{ESS1}	P_{min}	P_{max}	X_S	X_T	I_{f_0}	$I_{f_{max}}$
20 kWh	1 kWh	19 kWh	2 kWh	98%	97%	0.2 pu	1 pu	0.1 pu	1 pu	1 A	3.6 A

Fig. 8. PQt profile and belief function for UL agent.

4. Simulation results

Below we present the comparison between the studied control strategies, which is followed by the validation of the employed aggregation methods.

4.1. Short-term behaviour

In this section, we compare the results obtained in the scenario described in Section 2.3, with three different control methods: Commelec, Droop with primary frequency and voltage control only in the slack resource (DP), and Droop with additional secondary frequency control (DPS) at the slack. The focus here is on the *dynamic*

short-term behaviour. In particular, the results are presented over the time horizon of 1600 s.

4.1.1. Control of the reserve of the storage systems

The evolution of the state of charge (*SoC*) of both battery systems is shown in Fig. 9. Note that in the case of Commelec, the *SoC* decreases towards the target value (*SoC* = 0.5) as opposed to DP/DPS, in both LV and MV networks. In the case of the LV battery, when using Commelec, the *SoC* decreases much faster because this resource is being requested to discharge mostly at full power; whereas in the case of the MV battery, it is discharging but subject to the fact that this resource is the slack bus of the system (therefore its power production/absorption is the result of all other resources).

The evolution of the *SoC* of the WBs is also presented in Fig. 9. It can be seen that the boilers are controlled to react to local power variations while following their willing to be charged. WB1 and WB3 are being charged from the beginning at full power, whereas WB2 is charged when possible. On the contrary, in DP/DPS, they are not charged at all.

4.1.2. Reduced curtailment of renewables

Fig. 10 shows the production of the PVs, by means of the PV active power and the total produced PV energy. It can be seen that in Commelec, the PVs produce at maximum available power most of the time; whereas in DP/DPS, they are curtailed given the excess of power in the network assessed by the frequency signal. In this respect, with Commelec the renewables production is maximized even with high variability profiles and it is curtailed only when it
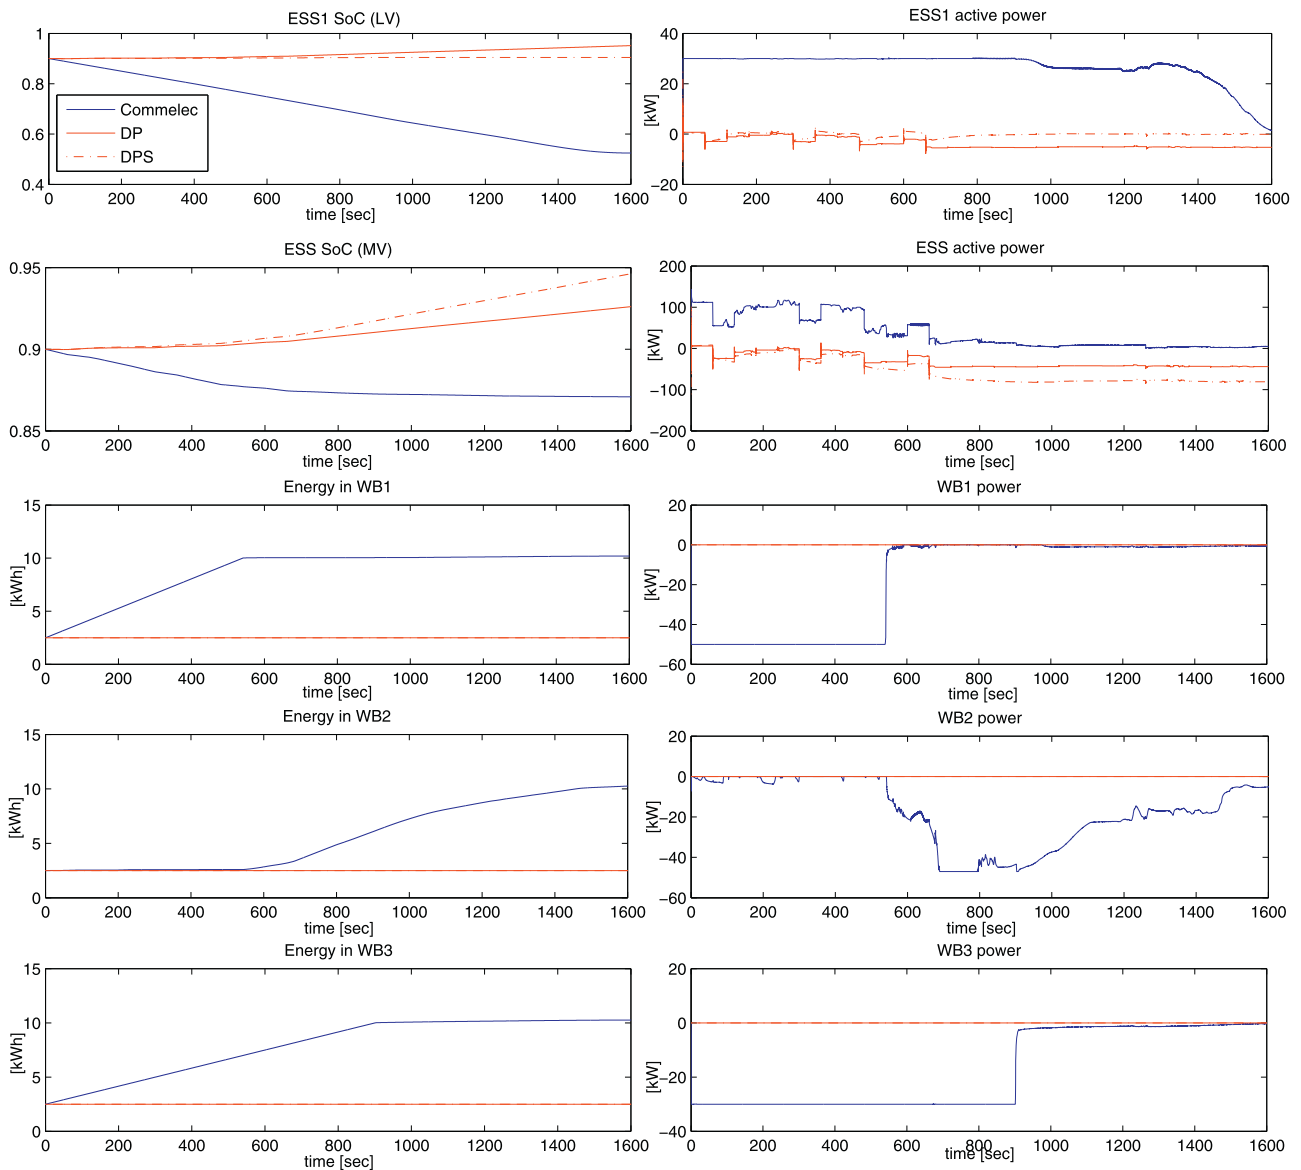
Fig. 9. Comparison between Commelec and both droop strategies. The left column presents the state-of-charge of the battery systems and the stored energy in the water boilers, while the power profile of the same elements is presented in the right column.

affects the power quality or there is not enough storing capacity in the system.

4.1.3. Local power compensations and exploitation of degrees of freedom

Fig. 11 shows the production of the synchronous generators (SG and μH). It is worth noting that in the case of Commelec, the power variations in the LV grid are compensated locally by the means of μH, while maintaining the MV SG at minimum power. In the droop simulation, on the contrary, both machines react in the same way. The main reason for this difference is that Commelec exports and use the internal state of the resources, whereas in DP/DPS the control is performed via the global frequency signal.

It is interesting to observe specifically the case of WB2, which is connected to the same bus with PV3 (see Fig. 1). This node is then connected to the main feeder of the LV network by a line with an ampacity close to the current being absorbed by WB2 at its rated power. We show the dynamic behaviour of these two devices in Fig. 12. It can be seen that WB2 starts charging around $t = 550$ s.

This becomes possible due to the overall state of the system, and in particular due to the fact that WB1 stops charging at this time (see Fig. 9). However, due to low production from PV3 at this time and the weakness of the line that connects both devices to the network, the charging is not at the maximum possible power. When the production of PV3 increases at around $t = 650$ s, WB2 starts charging at maximum power. We note that the line current remains below the ampacity during the whole process. This case illustrates again the ability of our method to compensate for power imbalances locally and to exploit the various degrees of freedom of the system by using the advertised information about the internal state of the devices.

4.1.4. Quality of service and stable frequency

In Fig. 13, the system frequency is presented. Recall that the slack is the MV storage system (ESS). As the Commelec method is explicit, the slack works at constant frequency.

In the case of DP, on the contrary, the frequency reacts to the changes in UL; whereas in the case of DPS, it tries to return to the reference value. It is important to note that the frequency variations
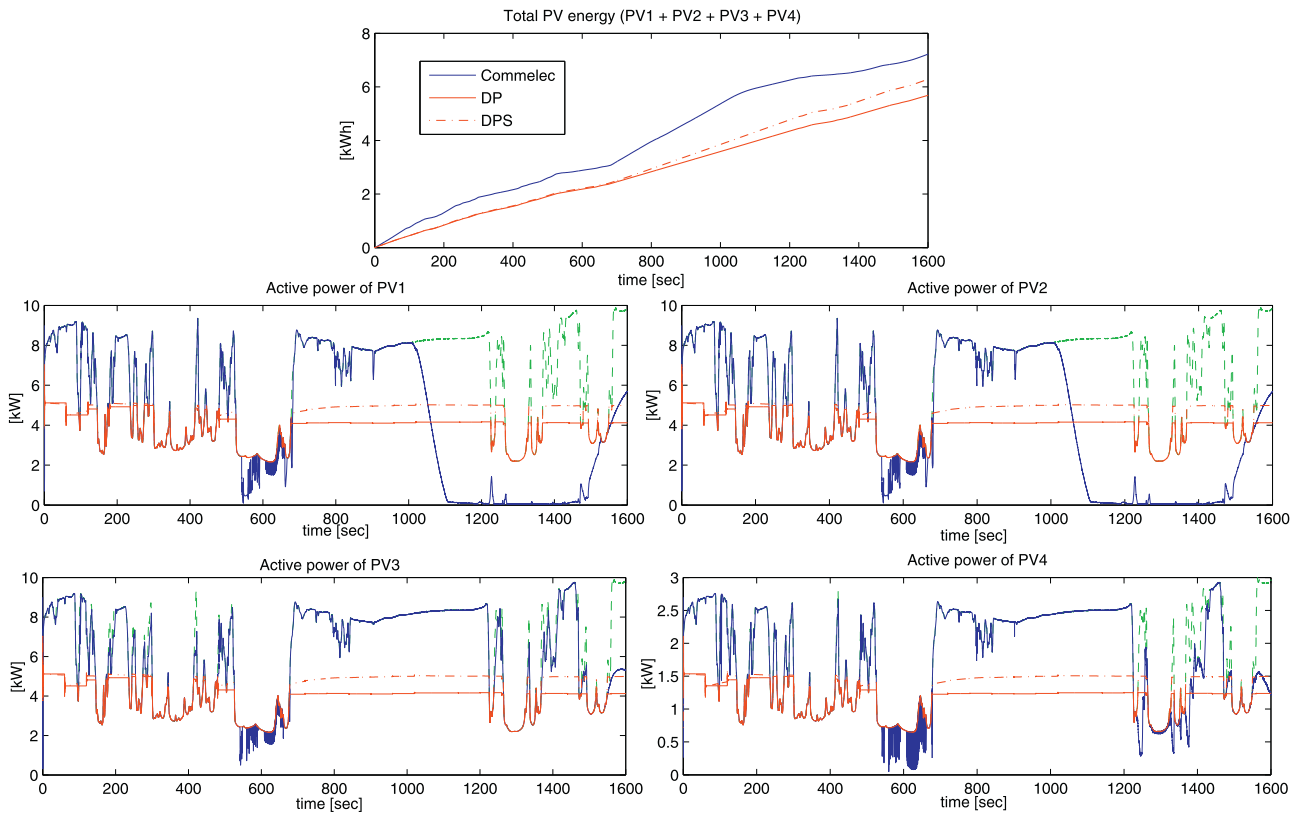
Fig. 10. Comparison between Commelec and both droop strategies. Total produced energy for the four PV plants connected to the LV microgrid and the power production for each. The dashed green line represents the maximum power production following directly the irradiance profile.

are highly dependent on the droop parameters and can be very high when there is a sudden change in the network. Therefore, by keeping the frequency constant, Commelec allows for a more accurate control of electrical machines. This is true especially in a microgrid that, when controlled using standard droop-based strategies, is expected to face high variations of frequency due to the uncertainty of the renewables.

In Fig. 14, we present the aggregated voltage and current profiles for both networks (i.e., median, minimum and maximum values of node voltages and line currents). It can be seen that the improvement in the overall operation obtained by using our method does not affect the quality of service. The voltage and current magnitudes are always within the acceptable regions. Note that the maximum LV current using Commelec is always close to the ampacity. This specific case is related to WB2 and PV3 as explained before. Observe,

however, that the median value is much lower during the entire simulation run.

4.2. Medium-term behaviour and system collapse

In this subsection, we illustrate the medium-term system behaviour in the critical case corresponding to the overproduction from renewables with initial high value of the *SoC* of the batteries and minimum load. Specifically, we present the *SoC*, the production of a PV, and the injection of SG and μH in Fig. 15, over the time horizon of 4000 s (namely, around 1 h). It can be seen that both DP and DPS control strategies lead to the overcharge of the MV battery, essentially causing the *collapse of the system*. In particular, when the power is injected into the ESS with *SoC* = 1, the local controller of the resource trips its breaker, with the consequent loss
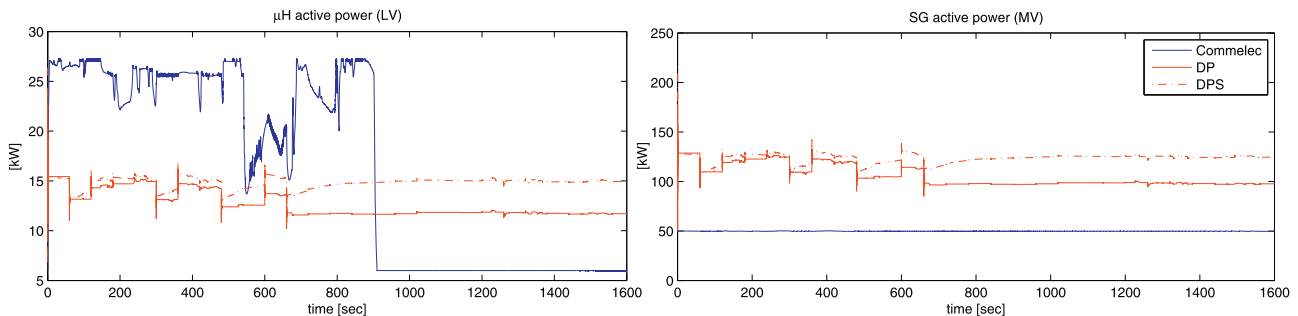

Fig. 11. Results for the comparison between Commelec and both droop strategies. Active power production of SG and μH.
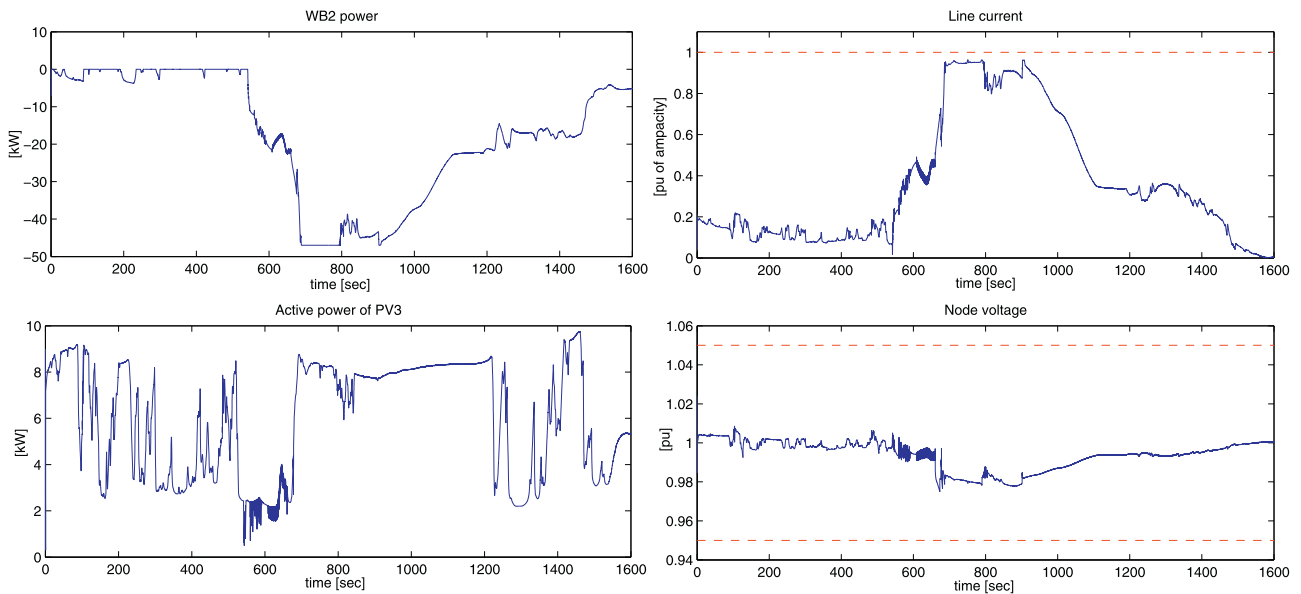
Fig. 12. Local power management between WB2 and PB3. The left column shows the power profiles while the right column shows the current of the line connected to both resources and the voltage of the common node. The dashed red lines represent the bounds.

of the slack resource provoking the collapse. The main reason for this behaviour is that the droop strategies force the generators to overproduce power regardless of the *SoC* of the slack resource. It is worth noting that in DP, as there is a permanent positive frequency error, the LV battery (ESS1) is always being charged. Hence, it trips even before the MV battery (ESS). The early loss of ESS1 can be also interpreted as a lack of autonomy of the microgrid if islanded. In the case of DPS, the secondary frequency control allows for a larger production of the generation units, and therefore the *SoC* of ESS1 is essentially constant. As a result, the MV battery is charged without restriction. On the contrary, Commelec keeps the *SoC* of both ESS and ESS1 away from the margins by using internal information from each resource and controlling explicitly their power setpoints.

4.3. Unexpected disconnection of a device

In this subsection, we demonstrate how Commelec is able to cope with an unexpected disconnection of a device. In particular, at $t = 1000$ s the resource PV1 and its agent are disconnected. Immediately, the slack resource (ESS) reacts to cope with the

imbalance. We remind that, as mentioned above, the control of the slack resource can be astatic. As a consequence, the compensation performed by the slack has no impact on the system frequency. Afterwards, Commelec takes over. Note that PV1 is directly connected to WB1 and PV2. As WB1 is already close to its minimum power, and PV2 aims at producing at maximum, the algorithm also reduces the consumption of WB2 and WB3 (connected to different nodes) to assist the maneuver. The simulation shows how Commelec handles unexpected disconnection by assisting the slack bus in redistributing the power imbalance between the resources and by keeping the overall state of the grid feasible (Fig. 16).

4.4. Validation of the aggregation methods

In this section, we numerically validate the aggregation methods described in Section 6 of Part I. To this end, we perform a simulation of the "flat" setting of agents shown in Fig. 2(b). In order to make a fair comparison between the results obtained in the hierarchical setting, Fig. 2(a), and those obtained in the flat setting, we adjust the weights of the objective function and the step-size parameters
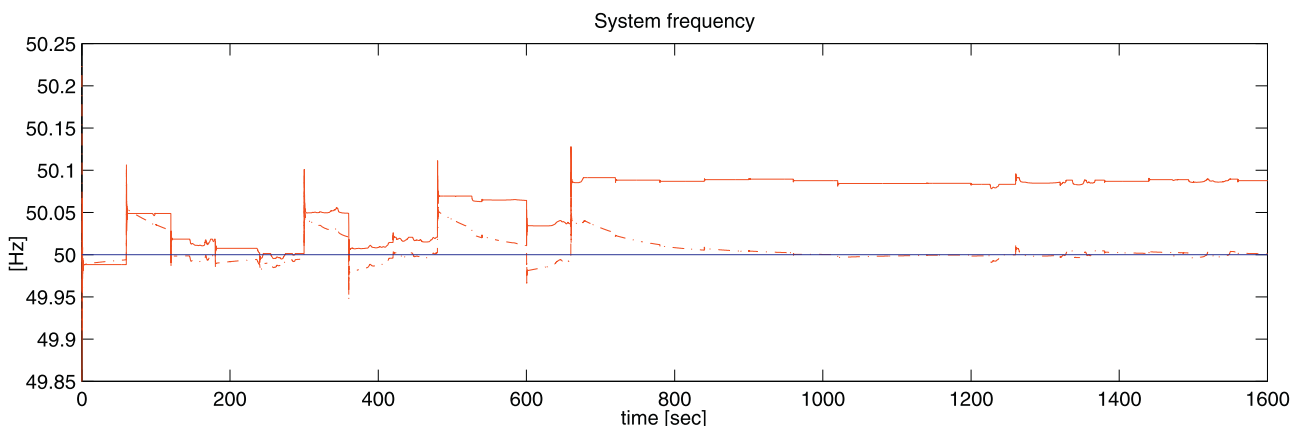

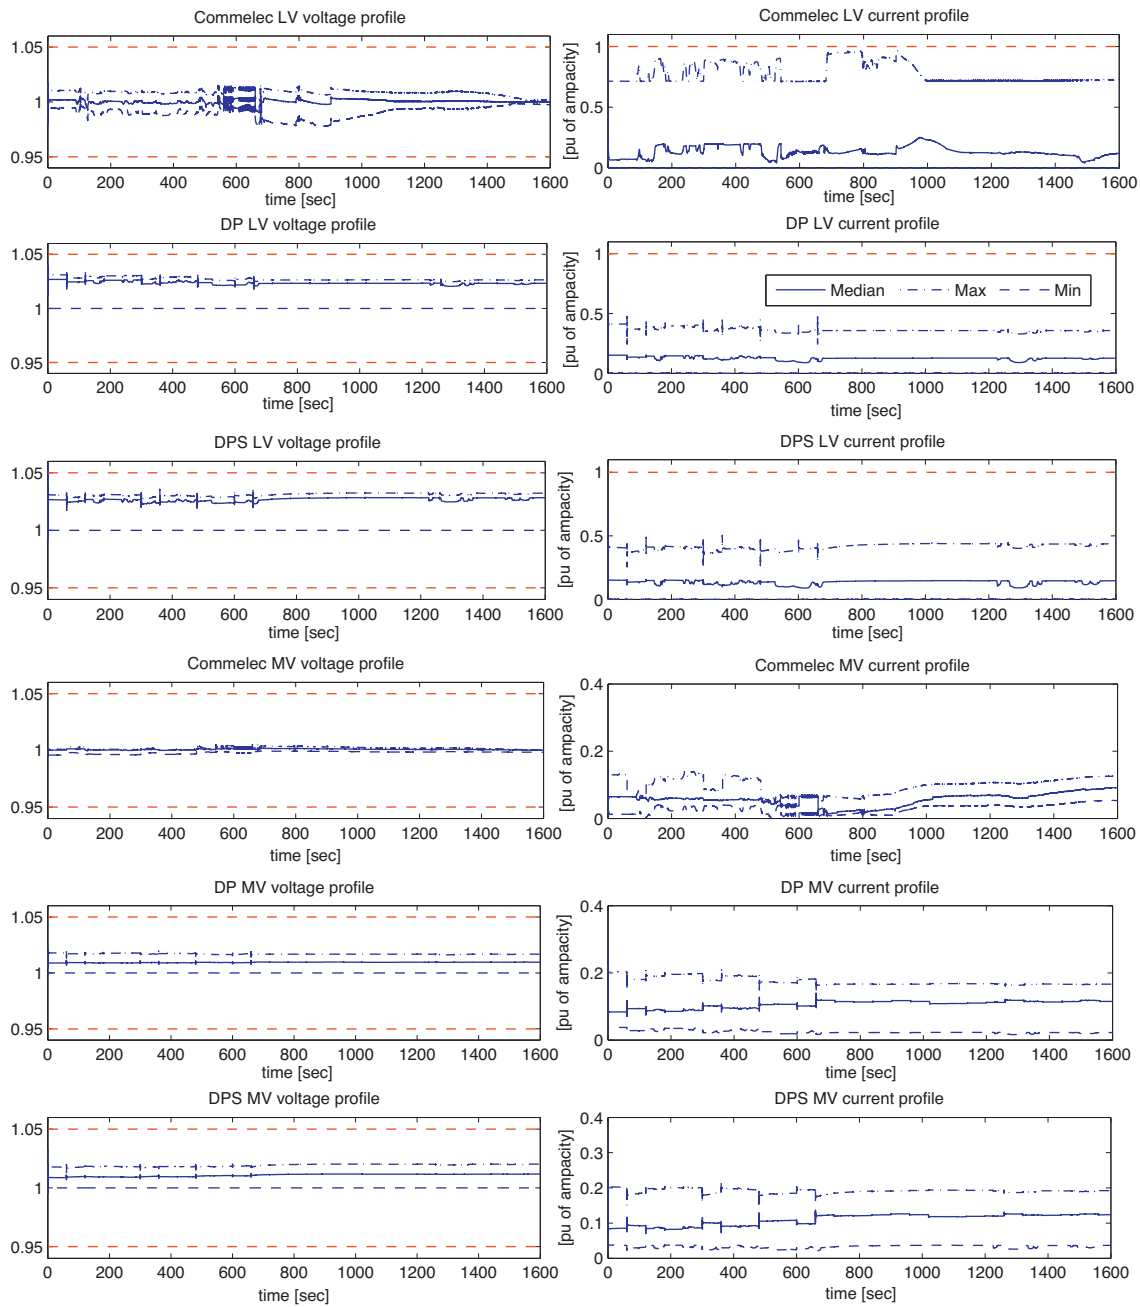**Fig. 13.** Comparison between Commelec and both droop strategies. System frequency.

Fig. 14. Comparison between Commelec and both droop strategies. The red dashed lines represent the predefined bounds for voltage and line congestions.

accordingly. In particular, the step-size parameter in the flat case is set to that of the MVGA, and the weights of the microgrid resources are multiplied by the ratio between the step-size of the LVGA and MVGA.

Fig. 17 presents a comparison between the results obtained in the two settings. As it can be observed, the overall behaviour is similar. The main difference is in the power of the synchronous generators, where a difference of up to 20 kW can be observed in the injection of SG. However, the contribution of this difference to the overall behaviour is negligible, as can be inferred from the presented energy metrics (*SoC*, and PV and WBs energy).

As shown in Proposition 5.1 in Part I, the two settings are equivalent under the "ideal" conditions stated there. In our implementation, however, there are three main reasons for the observed difference. First, there is a natural difference due to the approximate methods used for aggregation. Second, recall that we implement

a gradient-based algorithm rather than solve an exact optimization. Moreover, in the hierarchical setting, the LVGA is requested to provide a certain fixed power at the connection point, whereas in the flat setting this power can vary without any prescribed restrictions. Third, the projection algorithms used to compute the control are *randomized*. In particular, in Algorithm 2 presented in Appendix A in Part I, in order to efficiently find the direction of minimum violation, we draw setpoints uniformly and randomly.

5. Discussion

This section presents a general discussion of the proposed control framework in this two-part paper, with a focus on extensions and future directions.
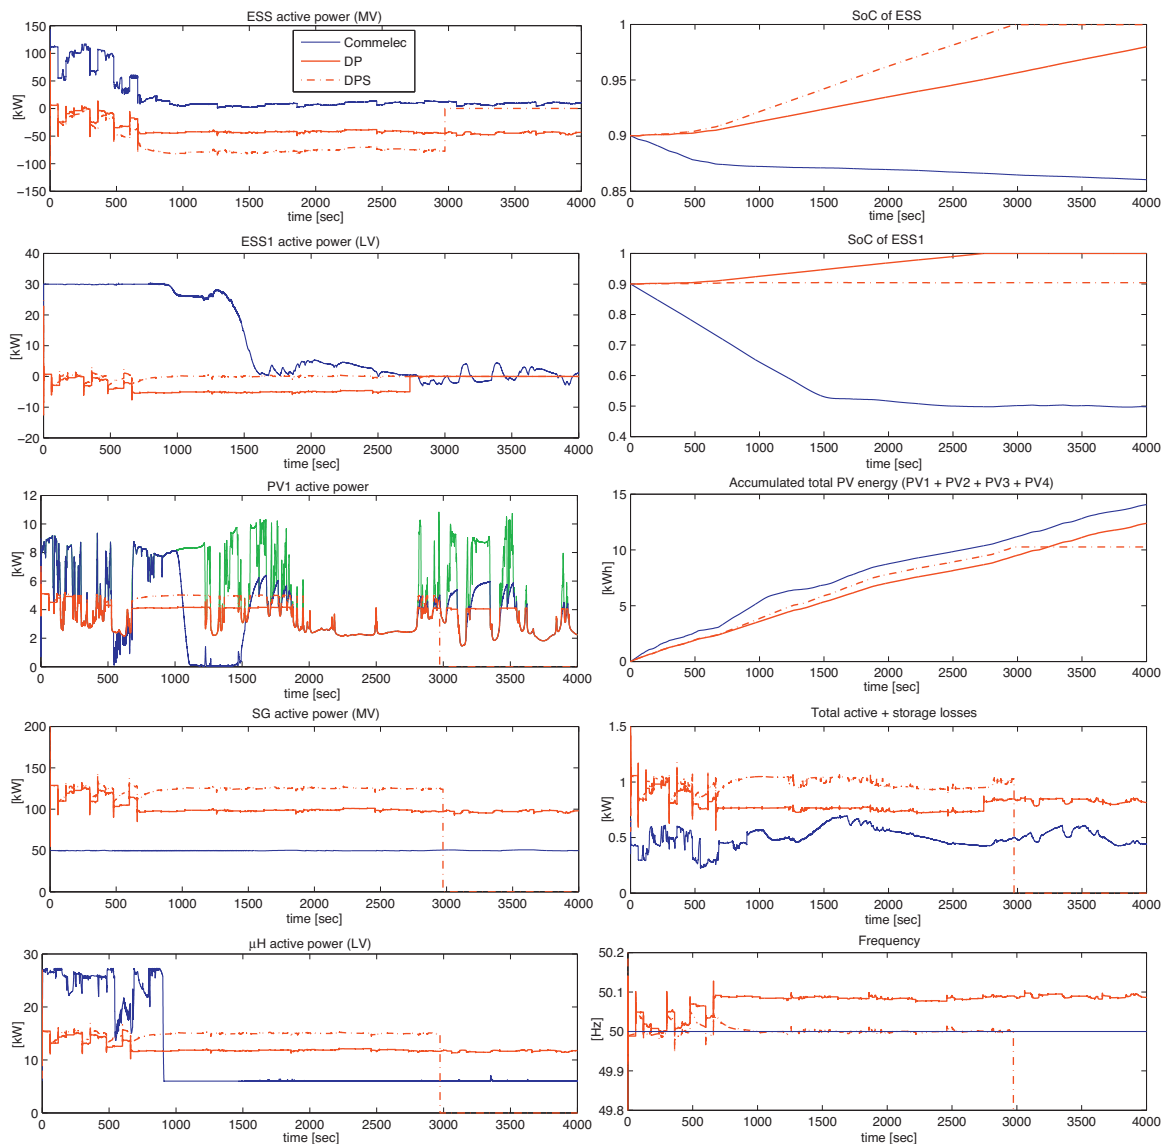
Fig. 15. Comparison between Commelec and both droop strategies. Medium-term comparison where the batteries are overcharged using DP/DPS strategies due to the production of renewables, even when curtailing their production.

5.1. Extension to higher-level grids

As shown in this part of the paper, the proposed framework achieves several desirable performance goals in an islanded distribution network, in the presence of highly volatile resources. These goals are achieved through a simple and generic method, with a key property of composability. It can be seen that this property allows us to easily extend our method to higher levels of power grids, up to the transmission level.

5.2. Partially controlled grids

In this paper, we assume that a grid agent faces resources that are fully controlled by Commelec agents. In a practical deployment, however, it is envisioned that certain parts of the grid cannot be controlled explicitly. It is possible to extend our framework to this "partially controlled" case. In particular, a "shadow agent" can be attached to each "non-Commelec" part of the grid. This shadow agent will take place of the regular Commelec agent. It will monitor the behaviour of that part, and estimate its model (by

using, e.g, network equivalence methods [11] or heuristic forecasting tools). By using this model, the agent can forecast the power production/consumption of that part of the grid and capture the uncertainty of this forecast in a belief function (similarly to the uncontrolled load agent), which in turn can be used in the optimization performed by the GA.

5.3. Robustness to faults and their treatment

We note that the proposed method naturally relies on the communication infrastructure for transmitting messages. In this paper, in the simulated case study, we assumed a perfect communication channel, with no message losses. Moreover, we did not take into account a possible failure of agents to produce advertisement messages, which might lead to incomplete information at the leader side. In this subsection, we outline how our method can be extended in order to make the communication between the agents more robust, and how to treat the failures when they occur.
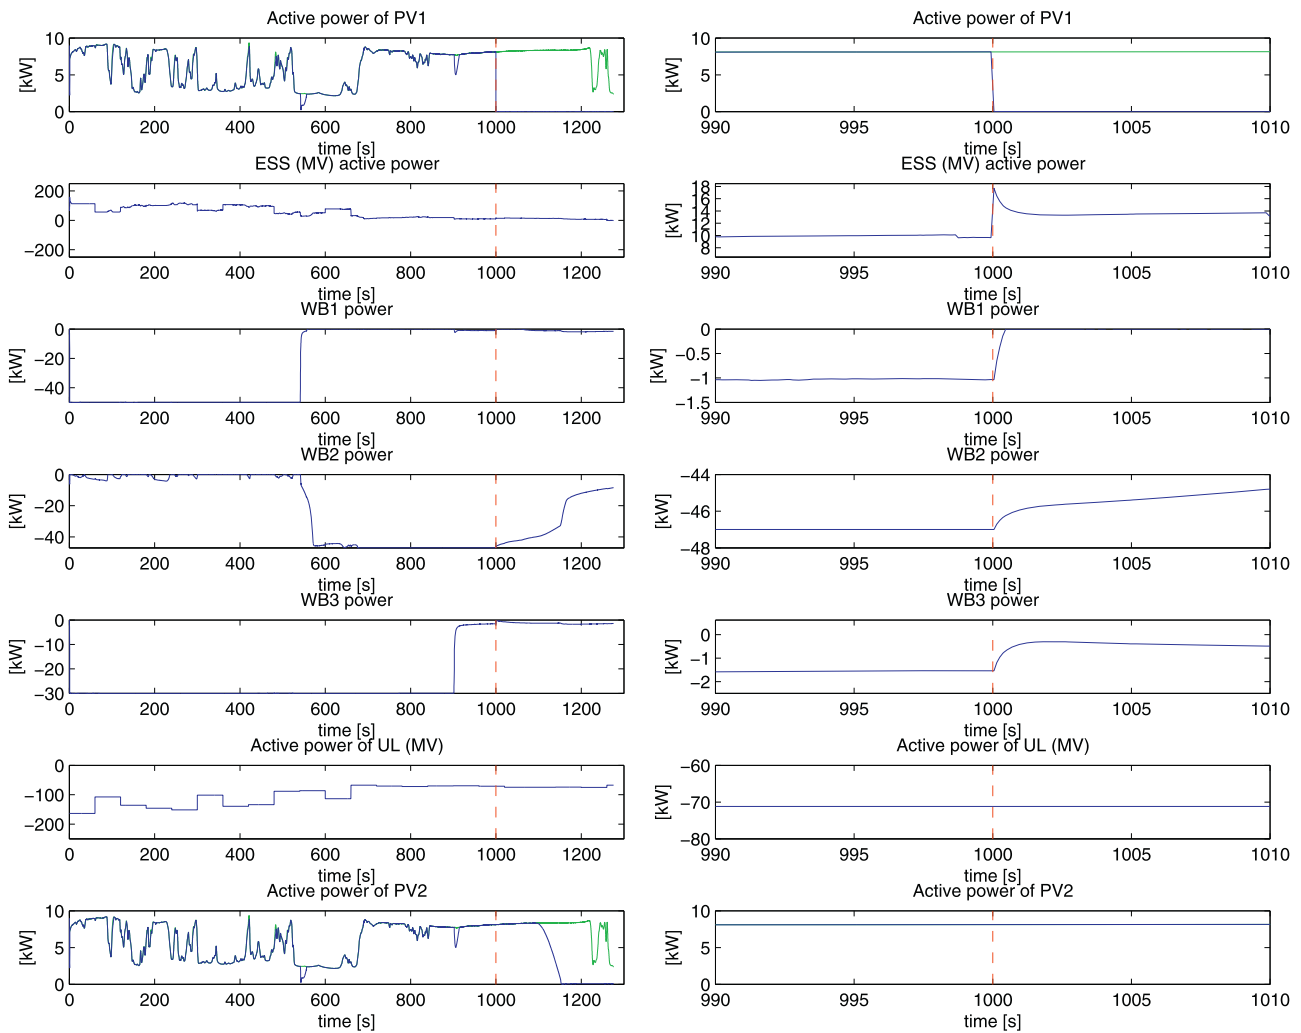
Fig. 16. Unexpected disconnection of PV1 at $t = 1000$ s. The right column shows a zoom on the left column in the time window of [990, 1010] s.

First, we envision a *triplicated implementation* of the grid agents (and possibly some important resource agents, e.g., those responsible for storage devices). In a normal mode of operation, each copy will perform the same computation and issue the same requests and advertisements. The triplicated data will be then sent to the receiver agent that will validate it using standard methods for validation of triplicated data (e.g., using voting mechanism based on some distance metric between the messages). Provided that the GA copies are synchronized to a certain extent, such a method would ensure correctness in the face of a failure of a single copy and/or loss of a single message. The communication protocol can be easily adapted to account for triplicated messages.

Second, a special communication networking infrastructure is envisioned to prevent packet losses; traffic engineering will be used to reduce congestion losses; source coding for long messages will mitigate the effect of packet losses; and a parallel redundancy protocol [12] will be used to provide instant packet loss repair.

Next we outline how failures can be treated when they occur. From the viewpoint of a follower agent, if there is no valid request setpoint from the leader (either due to the loss of messages or to validation failure), the agent can move to a *backup mode*. In this mode, the agent will produce setpoints according to some internal

decision process, within the feasibility of the system, with the little information it has. For instance, a RA can use a droop-based control, whereas a GA can operate in a similar way as described in this paper but without the term penalizing the deviation from the request at the connection point.

From the point of view of a leader agent, a "shadow agent" will be attached to each of its followers. If a failure is detected (e.g., if the advertisement message from a follower is not received for long time, or the message validation procedure fails), the shadow agent will take the place of the real agent. A shadow agent will have a functionality and goals similar to those responsible for the uncontrolled part of the grid discussed in Section 5.2.

5.4. Islanding maneuver and choice of local slack resources

In this paper, we focus on an islanded system in order to show that our method is able to operate a microgrid in an autonomous way, locally compensating for power imbalances. However, the proposed framework can be extended in order to allow for the *islanding maneuver* of a connected active distribution network. In particular, given a command from the leader to perform this maneuver, the grid agent will steer the system towards the state with 0 power at the connection point. At the same time, it will
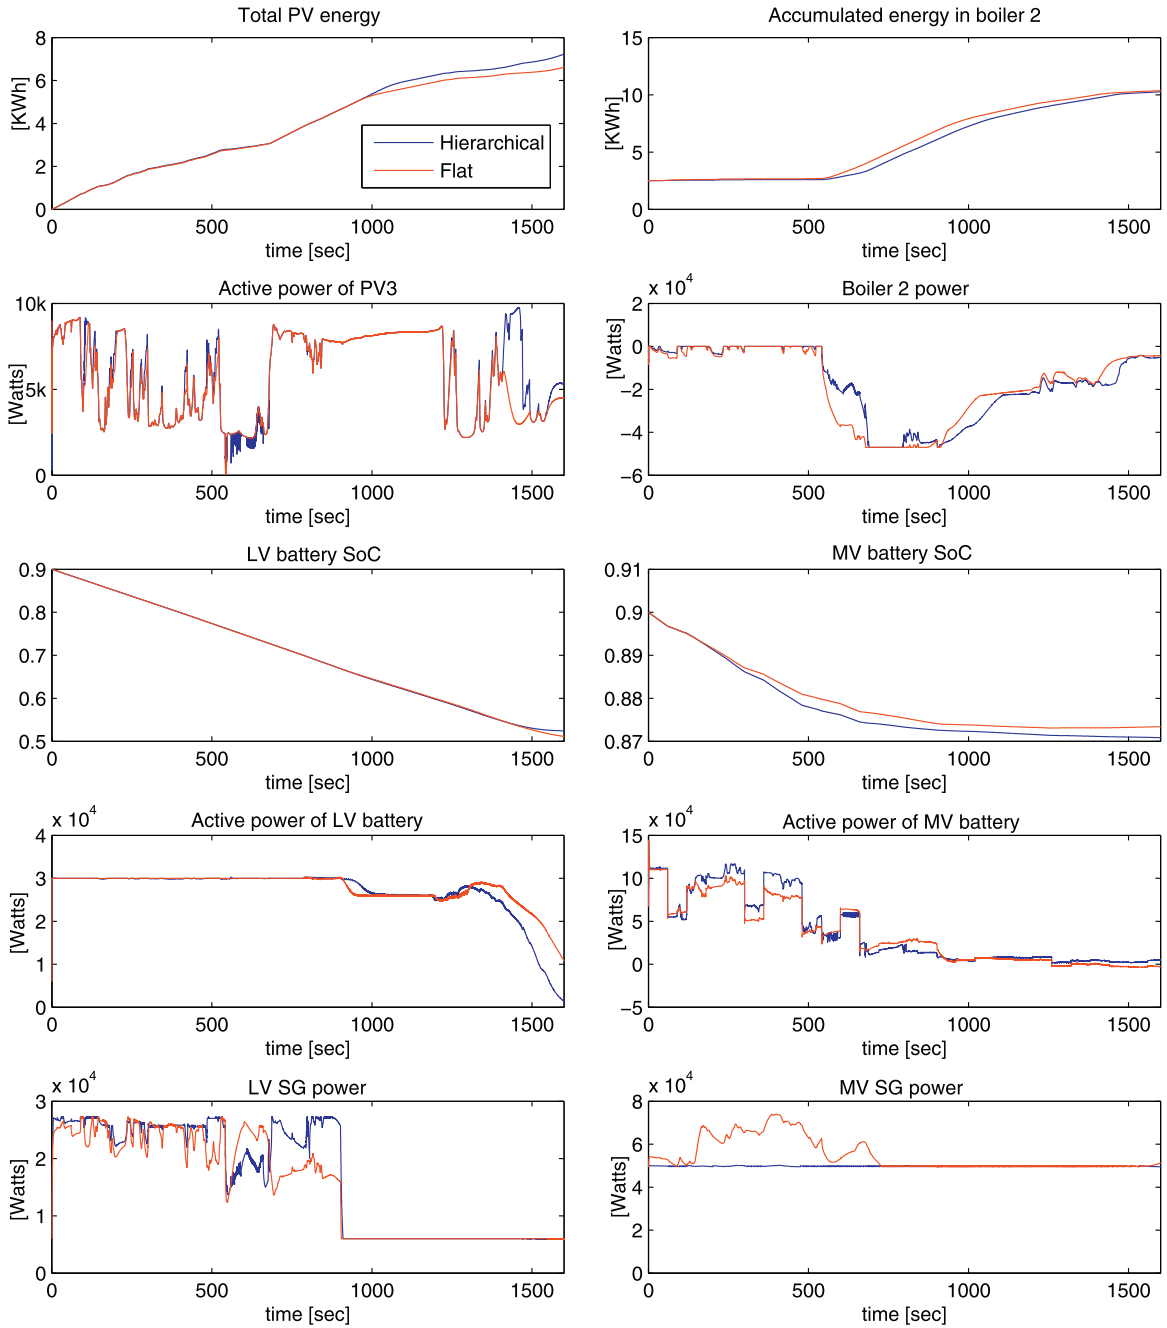
Fig. 17. Comparison between the hierarchical and flat agents architecture using the Commelec protocol.

perform a "negotiation" with its followers in order to choose *a set of slack resources*. We note that the grid agent can take its decision based solely on the advertised information from the followers (e.g., using a metric as in [13]). In particular, it will prefer to choose a resource with (i) a "good" belief function (e.g., battery or SG), (ii) a large range of available power as represented by the current *PQt* profile, and (iii) an internal state far away from the margins as represented by the advertised cost function.

Similarly, a *reconnection maneuver* can be implemented by steering the system towards a common frequency and voltage phasor at the connection point. In particular, we can add a further term in the GA objective function $J(y)$ accounting for the difference of the voltage phasors between the microgrid and the upper (larger) network, and we can instruct the slack resource in the microgrid to follow a common frequency.

5.5. Slack voltage control

In the current implementation, we assumed that the voltage at the slack bus of the system is fixed. Hence, it is not considered as a control variable. Moreover, the grid agent responsible for the slack does not have a way to decide which slack voltage is good or bad for the system, because it does not receive any related information from its followers. For instance, in our case study, the LVGA may prefer to increase the voltage due to high consumption in the microgrid, but the MVGA does not have a way to obtain this information. Still, we can easily adapt our framework to treat this case. Specifically, a follower agent can export an additional cost function to its leader, which gives a cost to specific value of the voltage magnitude at the connection point. Then, the leader that is responsible for the slack bus can incorporate these functions in the overall

optimization problem in order to choose an optimal voltage at the slack bus.

5.6. Incorporation of long-term objectives

We note that when considering resources equipped with storage systems (such as batteries, heating systems, etc.), the related long-term objectives can be incorporated easily in our framework by using the advertised cost functions. This can be achieved using a stand-alone "trip planner" (that is not necessarily part of the specific resource agent) that works on a slower time scale. For example, consider a trip planner responsible for controlling a residential building. Typically, it will have access to long-term forecasts of consumption and production patterns. It can compute an optimal control strategy by solving a multi-time step optimization problem using methods such as model predictive control (e.g., [14,15]). This computation is done usually on a time scale of tens of minutes. The trip planner can then "feed" the Commelec resource agent with a cost function that represents this long-term control strategy. Hence, the agent can advertise this information to its leader in order to be able to "steer" towards the trajectory prescribed by this strategy.

5.7. Probabilistic profiles

Recall that the PQt profiles sent by the followers are assumed to be deterministic sets in the PQ plane. As a result, a GA performs deterministic optimization under constraints imposed by the advertised belief functions. However, as the agents of volatile resources, such as PVs and loads, usually base their profiles on *forecasts*, probabilistic profiles can be considered. In particular, a probabilistic PQt profile can be considered as a collection of conditional probability distributions $p_t(x|u)$, with the interpretation that the actual setpoint x at time t is distributed according to $p_t(\cdot|u)$ whenever the control is u. Then, the grid agent will perform *stochastic optimization* using the expected value of the objective function with respect to the advertised distributions. In this case, the constraints using the belief function can be replaced by *chance constraints* using the distributions $p_t(x|u)$.

In this paper, however, we choose not to pursue this direction because the stochastic optimization framework poses several fundamental problems. It usually assumes independence of the underlying random variables, thus not taking into account the correlation between different resources (e.g., two PV farms located at the same geographical area). Moreover, normal distributions and noise independence are usually assumed to make the computation feasible. However, these assumptions do not necessarily hold in practice. On the contrary, developing a stochastic controller that accounts for the correlation between different resources as well as non-normal distribution of the volatilities requires the use of numerical approaches, thus computationally prohibitive for our real-time application.

6. Conclusion

We have introduced a method that uses explicit power setpoints in order to perform real-time control of electrical grids in a scalable and reliable way. The two main features of the proposed method are *correctness by construction* and *composability*. The applicability of the method was verified via simulations performed on a case study composed of a low voltage microgrid benchmark (proposed by Cigré Task Force C6.04.02) connected to a generic medium voltage feeder. The selected case study is characterized by (i) the typical level of complexity of distribution networks, (ii) a

pervasive penetration of renewable energy resources, (iii) the presence of distributed storage systems, and (iv) the fact that most of the inertia comes from storage and thermal loads rather than rotating machines.

The results of the performed simulations suggest that the proposed real-time control framework is able to efficiently steer such a system in the presence of extremely volatile energy resources. In particular, our findings show that (a) the method is able to indirectly control the reserve of the storage systems, thus maximizing the autonomy of the islanding operation, (b) it dramatically reduces the curtailment of renewables and is able to implicitly identify local power compensation, (c) it keeps the system in feasible operation conditions while better exploring the various degrees of freedom that characterize the system, and (d) that it maintains the system's power equilibrium without using he frequency as a global variable being able to do so in inertia-less systems. Most importantly, it prevents the system collapse in the case of over-production from renewables. We have also proven, by simulations, that the composability property of the proposed method holds. This specific peculiarity will potentially enable its application to generic and more complex power systems and further research efforts are expected in this respect. It can be concluded that the proposed control scheme represents an effective *actuation method* for the sub-second control of active distribution networks capable of accounting for the main requirements associated with the evolution of these grids.

References

[1] S. Papathanassiou, N. Hatziargyriou, 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", Athens, Greece, April, 2005.

[2] N. Hatziargyriou (Ed.), Microgrids: Architectures and Control, John Wiley and Sons Ltd, 2014.

[3] J. Matas, M. Castilla, L.G. de Vicu na, J. Miret, J.C. Vásquez, Virtual impedance loop for droop-controlled single-phase parallel inverters using a second-order general-integrator scheme, IEEE Trans. Power Electron. 25 (12) (2010) 2993–3002.

[4] F. Katiraei, M.R. Iravani, Power management strategies for a microgrid with multiple distributed generation units, IEEE Trans. Power Syst. 21 (4) (2006) 1821–1831.

[5] J. Grainger, W. Stevenson, Power System Analysis, McGraw-Hill, 1994.

[6] T. Reddy, Linden's Handbook of Batteries, 4th ed., McGraw-Hill Education, 2010.

[7] A. Rahmoun, H. Biechl, Modelling of Li-ion batteries using equivalent circuit diagrams, PRZEGLAD ELEKTROTECHNICZNY (Electr. Rev.) 88 (7) (2013) 152–156.

[8] B. Balvedere, M. Bianchi, A. Borghetti, C. Nucci, M. Paolone, A. Peretto, A microcontroller-based power management system for standalone microgrids with hybrid power supply, IEEE Trans. Sustain. Energy 3 (3) (2012) 422–431.

[9] K. Christakou, D.-C. Tomozei, J.-Y. Le Boudec, M. Paolone, GECN: primary voltage control for active distribution networks via real-time demand-response, IEEE Trans. Smart Grid 5 (2) (2013) 622–631.

[10] G. Costanzo, A. Bernstein, L. Reyes Chamorro, H. Bindner, J.-Y. Leboudec, M. Paolone, Electric space heating scheduling for real-time explicit power control in active distribution networks, in: Proceedings of the 3rd IEEE PES Innovative Smart Grid Technologies (ISGT) Europe Conference, Istanbul, Turkey, October 12–15, 2014.

[11] S. Zali, J. Milanovic, Generic model of active distribution network for large power system stability studies, IEEE Trans. Power Syst. 28 (3) (2013) 3126–3133.

[12] M. Popovic, M.M. Maaz, D.-C. Tomozei, J.-Y. Le Boudec, IPRP: Parallel Redundancy Protocol for IPv6 Networks, 2014 http://infoscience.epfl.ch/record/199448, Tech. rep.

[13] A. Borghetti, M. Bosetti, C. Nucci, M. Paolone, Dispersed energy resources scheduling for the intentional islanding operation of distribution systems, in: Proceedings of the 16th Power Systems Computation Conference (PSCC '08), Glasgow, Scotland, July 14–18, 2008.

[14] R. Halvgaard, N. Poulsen, H. Madsen, J. Jorgensen, Economic model predictive control for building climate control in a smart grid, in: Proceedings of the 2012 IEEE PES Innovative Smart Grid Technologies (ISGT), January 16–20, 2012.

[15] F. Oldewurtel, A. Parisio, C.N. Jones, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, M. Morari, Use of model predictive control and weather forecasts for energy efficient building climate control, Energy Build. 45 (2012) 15–27.