

Aggregation of Power Capabilities of Heterogeneous Resources for Real-Time Control of Power Grids

Andrey Bernstein*, Jean-Yves Le Boudec*, Mario Paolone[†], Lorenzo Reyes-Chamorro[†] and Wajeb Saab[‡]

* Laboratory for Communications and Applications 2

École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland
andrey.bernstein@epfl.ch, jean-yves.leboudec@epfl.ch

[†] Distributed Electrical Systems Laboratory

École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland
mario.paolone@epfl.ch, lorenzo.reyes@epfl.ch

[‡] Rigorous System Design Laboratory

École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland
wajeb.saab@epfl.ch

Abstract—Aggregation of electric resources is a fundamental function for the operation of power grids at different time scales. In the context of a recently proposed framework for the real-time control of microgrids with explicit power setpoints, we define and formally specify an aggregation method that explicitly accounts for delays and message asynchronism. The method allows to abstract the details of resources using high-level concepts that are device and grid-independent. We demonstrate the application of the method to a Cigre benchmark with heterogeneous and low-inertia resources.

Index Terms—Microgrids, aggregation, real-time control, explicit control, heterogeneous energy resources

I. INTRODUCTION

The aggregation of heterogeneous energy conversion systems connected to medium and low voltage distribution grids is a fundamental functionality required in order to quantify the level of flexibility that distribution grids can offer to the bulk network. It is typically used as an input for control and dispatch functions operated at higher control levels.

The complexity of the aggregation depends on the number of the aggregated resources (i.e., generators, storage systems and loads), on the time scale where this function is performed (i.e., day-ahead, minutes or subseconds) and on the approach used in the aggregation process. Concerning this last element, the literature has proposed two main approaches. The first one relies on the use of a-priori defined models of a given family of homogeneous resources which parameters are identified via representative devices. The behaviour of these representative devices allows to capture the behaviour of the entire set of the aggregated ones (e.g., [1], [2], [3], [4]). The second approach is model-free since it infers the behaviour of the distributed devices from the interaction between them and a central unit (i.e., the aggregator) (e.g., [5], [6]). Usually, these approaches adopt data-driven learning techniques.

As mentioned above, the time scale of the aggregation can vary from day-ahead to subsecond as a function of the target application. In this respect, the majority of the literature has concentrated the efforts to define aggregation processes for electricity balancing markets with consequent aggregation time scales ranging from 24 hours to 15 minutes (e.g., [2], [7], [8], [9]).

The aggregation of heterogeneous resources for hard real-time applications (i.e., in the subsecond time scale) has been proposed in [10], [11], where the authors have proposed a solution for controlling a distribution network in real time using explicit power set-points. In this framework, the resources can advertise their current internal needs and power availability by simple messages in order to allow a grid controller to take wise and informed decisions and maintain the state of the system within secure limits. The framework, called *Commelec*, is designed to be robust (i.e., to avoid the problems inherently posed by software controllers) and scalable (i.e., it easily adapts to grids of any size and complexity). It is based on software agents, which are responsible for resources/subsystems (Resource Agents - RA) or entire grids (Grid Agents - GA) and communicate using a simple, yet powerful protocol with a refresh rate of around 100 ms.

One of the main characteristic of the *Commelec* framework is the possibility for a generic resource to advertise its internal state via common quantifiers representative of its capability limits and of its stochastic behaviour. This peculiarity allows a real-time and model-free aggregation of the resources via the quantification of their power capabilities, costs and stochastic behaviours. Another peculiarity of the framework is its applicability to systems with little or no inertia. Hence, the resulted aggregation provides higher flexibility than traditional approaches based on the assessment of aggregated inertia.

Concerning the aggregation, the *Commelec* control framework has been formulated with these two hypotheses on the information exchange between RAs and the GA: (i) no delays and (ii) synchronous exchange. Moreover, the proposed aggregation method in [10] was based on *sampling* of representative

Supported by the SNSF - NRP 70 "Energy Turnaround".

number of setpoints. The computational burden of this method might become prohibitively large for real-time application when the number of sample points grows. In this paper, we first propose a method that takes into account the delays and enables the asynchronous message exchange. The method is not based on sampling and is computationally efficient. Then, we show the application of the proposed method to aggregate largely heterogeneous devices in hard real-time (i.e., 100 ms) with respect to an application example referring to the microgrid benchmark defined by the CIGRÉ Task Force C6.04.02 which nodal injections/absorptions are defined from experimentally inferred time series measured on the EPFL campus.

The structure of the paper is the following. Section II summarizes the general characteristics of the Commelec control framework. Section III describes the formal aggregation method with particular reference to the time delays and synchronism aspects. Section IV illustrates the application of the Commelec aggregation on the benchmark microgrid. Section V concludes the paper by summarising the main features of the Commelec aggregation.

II. THE COMMELEC FRAMEWORK

An example of the Commelec agents structure is shown in Figure 1 (b), where the GA at the LV level (LVGA) is in charge of controlling a group of RAs responsible for specific subsystems, while the GA at the MV level (MVGA) is responsible for controlling LVGA and three additional RAs. The agents relation corresponds to the case study shown in Figure 1 (a), which is used here to illustrate the performance of the proposed methods – see Section IV for further details.

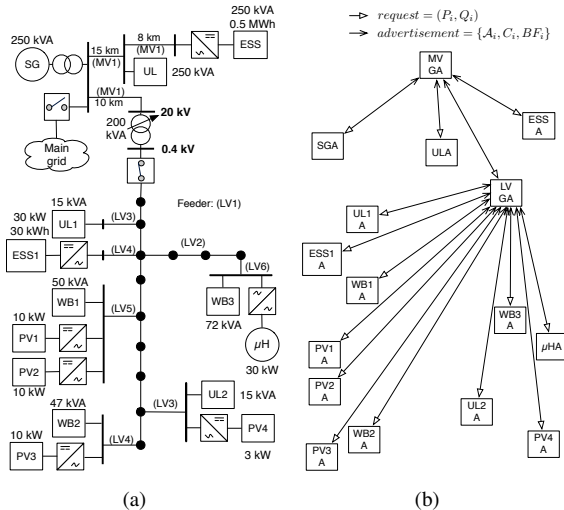


Figure 1. The electrical network and agents for the case study. (a) Electrical grid. (b) Agents. PV: Photovoltaic plant. UL: Uncontrollable Load. WB: Water Boiler. ESS: Energy Storage System. μ H: Micro-Hydraulic generator. SG: Synchronous Generator. LVGA: Low-Voltage Grid Agent. MVGA: Medium-Voltage Grid Agent.

Each GA is assigned a role of a *leader* of one or more other agents that we term the *followers* of that leader. The roles follow the hierarchy of distribution and transmission networks.

In the example of Figure 1, the LVGA is a leader of the RAs in its grid, while is a follower of the MVGA.

The framework uses a common, device-independent protocol for message exchange between the agents. In particular, each follower agent advertises its internal state to the leader GA using the following three elements.

PQ Profile and Virtual Cost. The *PQ profile* $\mathcal{A}^\ell \subseteq \mathbb{R}^2$ of follower $\ell = 1, \dots, L$ is the region in the (P, Q) plane (for active and reactive power) that the subsystem under the control of this agent can deploy (*negative power means consumption*). The *virtual cost* function, $C^\ell : \mathcal{A}^\ell \rightarrow \mathbb{R}$, is interpreted as the cost to this subsystem of applying a requested power setpoint. Its role is to quantify the propensity of this subsystem to deploy (P, Q) setpoints. Note that the cost is virtual and does not represent money.

The *overall PQ profile* is given by the Cartesian product $\mathcal{A} = \mathcal{A}^1 \times \dots \times \mathcal{A}^L$.

Belief Function. The *belief function* BF^ℓ returns the set of all possible (actual) setpoints that the subsystem under the agent control *might* implement. Specifically, assuming that the subsystem receives from its leader GA a request to implement a setpoint (P, Q) , the *actual* setpoint (P', Q') that this subsystem does implement lies in the set $BF^\ell(P, Q)$ with overwhelming probability. The belief function accounts for the uncertainty in subsystem operation. In particular, highly controllable subsystems, such as batteries and generators, are expected to have (almost) ideal beliefs, namely $BF^\ell(P, Q) = \{(P, Q)\}$. For subsystems such as PV/wind farms, or loads, the belief function will return larger sets, to account for their volatility.

The *overall belief function* is defined by the Cartesian product $BF(P^1, Q^1, \dots, P^L, Q^L) = BF^1(P^1, Q^1) \times \dots \times BF^L(P^L, Q^L)$.

At every time step, a given GA receives the following information: (i) the advertisements from its followers (with *PQ* profiles, virtual costs, and belief functions), (ii) the power setpoint request from a higher level (leader) GA, and (iii) the estimation of the electrical state of the grid (using real-time methods as in [12], [13]). There are two tasks that the GA performs. The first task is to steer the electrical state of its grid by explicitly setting the power setpoints so that (i) the virtual costs of its followers are minimized, (ii) the request from a leader is satisfied as much as possible, and that (iii) the grid is in a *feasible state of operation*. The latter refers to *static* (rather than dynamic) feasibility, defined in terms of the nodal voltage magnitudes and line currents, as in [10]. We note that this static analysis is reasonable as we focus on microgrids with little or no inertia, with resources connected to the grid by power electronic interfaces.

The second task is to compute the *aggregated advertisement* that will be sent to the leader GA. The advertisement should reflect the actual flexibility and uncertainty of the grid under the control of the GA based on the advertisements received from the followers. As mentioned, in this paper we focus on this second task.

III. PROPOSED AGGREGATION METHOD

Consider, without loss of generality, a Commelec setting with three layers: RA, LVGA, and MVGA (see Figure 1 for an example). In this paper, we propose a method to perform aggregation on the LVGA layer that (i) allows for composability in the Commelec framework, (ii) supports “soft state approach”, namely does not require synchronized operation and is inherently robust to delays in communication, and (iii) always keeps the system in feasible electrical state (at all the layers).

We introduce some notation. Let $u_i^0 = (P_i^0, Q_i^0)$ be the i -th request received by LVGA from its leader MVGA and ADV_j be the j -th advertisement received by LVGA from its L followers. In particular, $ADV_j = (A_j, C_j, BF_j)$ contains the overall PQ -profile, belief and cost functions obtained from the followers. Similarly, let $u_j = (u_j^\ell)$ denote the j -th setpoint sent by the LVGA to its followers, which is the collection of the individual PQ -setpoints $u_j^\ell = (P_j^\ell, Q_j^\ell)$, $\ell = 1, \dots, L$, and $AADV_i = (AA_i, AC_i, ABF_i)$ denote the i -th advertisement sent by the LVGA to MVGA. We usually omit the indices i and j when their value is clear from the context.

Note that from the point of view of the LVGA, two elements have to be computed periodically: the setpoint to the followers u_j and the (aggregated) advertisement $AADV_i$. We first specify the *timing* of this computation and its *dependency* on the elements received by the LVGA, namely¹ u_{i-1}^0 and ADV_j . We then proceed to devise specific procedure for computation of u_j and $AADV_i$.

A. Timing and Dependency

The timing and dependency of the computation of u_j and $AADV_i$ is specified in Algorithm 1. A time diagram that illustrates this algorithm is shown in Figure 2.

Algorithm 1 Timing and Dependence of Aggregation

- **Initialize:** $i = 1, j = 1$.
 - Repeat forever:
 - 1) **Wait for obtaining either u_i^0 or ADV_j .**
 - 2) **On obtaining u_i^0 :**
 - a) $i := i + 1$
 - b) Compute and send the aggregated advertisement $AADV_i$ (“compAggr” in Figure 2) based on the information in ADV_j so that the aggregated belief function correctly reflects the uncertainty introduced by the last setpoint u_j sent to the followers. This is described in Section III-C.
 - 3) **On obtaining ADV_j :**
 - a) Compute and send the setpoint u_j to the followers (“compSP” in Figure 2) based on the information in ADV_j and the last request u_{i-1}^0 , so that *the uncertainty of the power delivered at the PCC to MVGA is compatible with the recently sent aggregated advertisement $AADV_i$* . This is described in Section III-B.
 - b) $j := j + 1$.
-

¹Implicitly, we assume the all the computations are done based on the latest information obtained from the state estimation procedure.

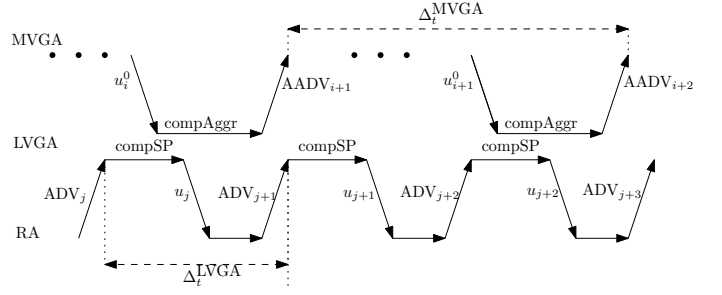


Figure 2. Timing illustration for aggregation and setpoint computation.

B. Computation of u_j

The computation is done similarly to [10], using a gradient steering algorithm:

$$u_j = \mathcal{P}_{\mathcal{U}_j} \{ \hat{x} - \alpha \nabla F_j(\hat{x}) \}, \quad (1)$$

where \hat{x} is the estimation of the actually implemented overall power setpoint, $\nabla F_j(\cdot)$ is computed from $ADV_j = (A_j, C_j, BF_j)$ and u_{i-1}^0 as in [10], and $\mathcal{P}_{\mathcal{U}_j} \{ \cdot \}$ is the projection operator onto the *admissible set* \mathcal{U}_j . The only difference from [10] is in the definition of the latter, that is modified to²

$$\mathcal{U}_j = \{ u \in A_j : \forall x \in BF_j(u), J(x) < \infty, X_0(x) \in ABF_i \}, \quad (2)$$

where x is the actually implemented overall power setpoint, $J(x)$ is the grid operation cost, $X_0(x)$ is the power at the PCC with MVGA when x is the power setpoint, and ABF_i is the recently computed aggregated belief function. In words, a setpoint u is said to be *admissible* if, no matter what the actual implementation is, the grid is in feasible electrical state ($J(x) < \infty$) and the power flow at the PCC is within the aggregated belief set advertised to the leader (MVGA).

C. Computation of $AADV_i$

There are many ways to perform aggregation of the advertisements from the followers. Here, our goal is to propose a *computationally efficient* method that ensures *feasibility of the electrical state* of the entire system – see Theorem III.1 below.

1) *Aggregated Belief Function:* In this paper, we assume that the ABF does not depend on the request u^0 , and we identify the corresponding constant function with the uncertainty set. This assumption is reasonable whenever the MVGA works at much slower time pace than LVGA. In addition, this assumption is reasonable since the control framework updates u_j with a refresh rate of 10Hz, much higher than the phasor dynamics of the distributed energy resources.

In principle, our goal is to expose to the MVGA the uncertainty at the PCC that is caused by the last implemented

²We note that this set is similar to the one computed in the islanded case (namely, for the MVGA in our case study), when the power at the PCC has to satisfy the constraints of the slack resource. For more details, see Section 5.2.2 in [10].

setpoint, namely u_j . As in [10], we can solve the following OPFs to obtain the rectangular superset:

$$\begin{aligned} & \max / \min P_0 && \max / \min Q_0 \\ \text{s.t. } & \begin{cases} x \in BF(u_j), \\ (P_0, Q_0) = X_0(x), \end{cases} && \text{s.t. } \begin{cases} x \in BF(u_j), \\ (P_0, Q_0) = X_0(x). \end{cases} \end{aligned}$$

Moreover, for computational simplicity, these problems can be further approximated by using bounds on power losses as in [10]. Let us denote the resulting rectangular set by $D(u_j) \subseteq \mathbb{R}^2$. The direct way then to expose the uncertainty would be to advertise $ABF_i = D(u_j)$. However, this can be a too restricted set in view of the admissibility set \mathcal{U} defined in (2). Specifically, when the new advertisement ADV_{j+1} arrives, it is possible that the admissible set becomes empty as it is not possible to satisfy the constraint $\forall x \in BF(u), X_0(x) \in ABF_i$.

To solve this problem, we propose to inflate the belief set as follows. Consider a ball $\mathcal{B}(u_j, \rho)$ of setpoints with radius ρ centered on u_j . Sample K setpoints $u^{(k)} \in \mathcal{B}(u_j, \rho)$, $k = 1, \dots, K$, so that $u^{(k')} = u_j$ for some k' . Then advertise

$$ABF_i = \text{rect} \left[\bigcup_{k=1}^K D(u^{(k)}) \right],$$

where $\text{rect}(\cdot)$ is the ‘‘rectangular hull’’, namely the corresponding bounding box approximation. Clearly, this set contains the uncertainty caused by the currently implemented setpoint u_j , as well as that of implementing other setpoints that are at the distance at most ρ from u_j .

2) *Aggregated PQ Profile*: In this paper, to make the computations simple and compatible with our real-time application, we choose to aggregate the PQ profile by ignoring the losses in the load-flow problem and by considering only the recently advertised PQ profiles from the followers. Namely, we assume that the power at the PCC takes form of $u^0 = \sum_{\ell=1}^N u^\ell$, where $u^\ell \in \mathcal{A}_j^\ell$, and \mathcal{A}_j^ℓ is the PQ profile advertised in ADV_j by follower ℓ . Hence, the aggregated PQ profile is given by

$$AA_i = \sum_{\ell=1}^N \mathcal{A}_j^\ell,$$

where the summation is the Minkowski set summation. We note that the Minkowski summation can be computed efficiently when the sets \mathcal{A}_j^ℓ are approximated by convex polygons.

3) *Aggregated Cost Function*: Under the assumption of slow change of the request $u^0 = (P^0, Q^0)$ from the MVGA and due to frequent setpoint updates, we propose to advertise a *linear approximation* of the cost function in the form

$$AC_i(P^0, Q^0) = \lambda P^0 + \mu Q^0,$$

where

$$\lambda \triangleq \partial F_j^* / \partial P^0, \quad \mu \triangleq \partial F_j^* / \partial Q^0$$

are the partial derivatives of the objective function of the LVGA computed at the current setpoint $u_j \equiv u$. We can write this function explicitly as follows [10]:

$$F_j^*(P^0, Q^0) = \sum_{\ell} w_{\ell} C_{\ell}(u^{\ell}) + J(u) + J_0(u, P^0, Q^0),$$

and note that u depends on P^0, Q^0 through (1). Thus,

$$\begin{aligned} \frac{\partial F_j^*}{\partial P^0} &= \sum_{\ell} \left[w_{\ell} \left(\frac{\partial C_{\ell}}{\partial P^{\ell}} \frac{\partial P^{\ell}}{\partial P^0} + \frac{\partial C_{\ell}}{\partial Q^{\ell}} \frac{\partial Q^{\ell}}{\partial P^0} \right) \right. \\ &\quad \left. + \frac{\partial J}{\partial P^{\ell}} \frac{\partial P^{\ell}}{\partial P^0} + \frac{\partial J}{\partial Q^{\ell}} \frac{\partial Q^{\ell}}{\partial P^0} + \frac{\partial J_0}{\partial P^{\ell}} \frac{\partial P^{\ell}}{\partial P^0} + \frac{\partial J_0}{\partial Q^{\ell}} \frac{\partial Q^{\ell}}{\partial P^0} \right] + \frac{\partial J_0}{\partial P^0}, \end{aligned}$$

and similarly for $\partial F_j^* / \partial Q^0$. Therefore, we need to compute $\partial P^{\ell} / \partial P^0$, $\partial Q^{\ell} / \partial P^0$, $\partial P^{\ell} / \partial Q^0$, and $\partial Q^{\ell} / \partial Q^0$, for $\ell = 1, \dots, L$. (The other partial derivatives can be computed analytically.) As the dependence of u^{ℓ} on u^0 in (1) involves projection, we propose to approximate these derivatives by assuming that the projection is not active, namely

$$\begin{aligned} u &= \hat{x} - \alpha \nabla F_j(\hat{x}) \\ &= \hat{x} - \alpha \left(\sum_{\ell} w_{\ell} \nabla C_{\ell}(\hat{x}^{\ell}) + \nabla J(\hat{x}) + \nabla J_0(\hat{x}, P^0, Q^0) \right). \end{aligned}$$

Using the definition of the quadratic J_0 given in [10], we have that, for example:

$$\begin{aligned} \frac{\partial P^{\ell}}{\partial P^0} &= -\alpha \frac{\partial}{\partial P^0} \frac{\partial J_0}{\partial P^{\ell}} \\ &= -\alpha \frac{\partial}{\partial P^0} \left(-2(P_0 - P'_0(\hat{x})) S_{P^{\ell} \rightarrow P^0} \right. \\ &\quad \left. - 2(Q_0 - Q'_0(\hat{x})) S_{P^{\ell} \rightarrow Q^0} \right) \\ &= 2\alpha S_{P^{\ell} \rightarrow P^0}. \end{aligned}$$

where $(P'_0(\hat{x}), Q'_0(\hat{x})) = X_0(\hat{x})$ is the power at the slack bus (PCC) when the power setpoint is \hat{x} and $S_{P^{\ell} \rightarrow P^0} \triangleq \partial P'_0(\hat{x}) / \partial P^{\ell}$, $S_{P^{\ell} \rightarrow Q^0} \triangleq \partial Q'_0(\hat{x}) / \partial P^{\ell}$ are the *slack power sensitivity coefficients* computed using a computationally efficient method similar to, e.g., [14]. Similarly,

$$\begin{aligned} \partial Q^{\ell} / \partial P^0 &= 2\alpha S_{Q^{\ell} \rightarrow P^0}, \\ \partial P^{\ell} / \partial Q^0 &= 2\alpha S_{P^{\ell} \rightarrow Q^0}, \quad \partial Q^{\ell} / \partial Q^0 = 2\alpha S_{Q^{\ell} \rightarrow Q^0}. \end{aligned}$$

D. Feasibility of Electrical State

We next claim that, by construction, our aggregation methods ensure feasible electrical state of the overall system. To that end, consider any hierarchy of agents in the Commelec framework with M levels indexed by $m = 1, \dots, M$. Let Δ_t^m denote the cycle of computation at level m , which depends on the computation and communication delays at m (see Figure 2). We use the terminology of [10] to denote any grid agent that has a leader an *internal GA*, while the grid agent at the root of the hierarchy a *root GA*.

Theorem III.1. *Suppose that the internal GAs compute setpoints and aggregated advertisements as described in Sections III-A, III-B and III-C; and that the root GA compute setpoints*

as in [10]. Assume that the advertisements sent by the resource agents at any level m are valid in the sense that when a new setpoint is applied, it stays within the advertised belief set during the period of at least Δ_t^m . Then, in the asynchronous setting described in Section III-A, the electrical state of the overall grid is feasible for any values of $\{\Delta_t^m\}$.

Proof. Fix a level m . Note that if all the advertisements from the followers are valid in the above mentioned sense, then the electrical state will be feasible at this level. This is true since the setpoints are chosen from the admissible set (2). Since the advertisements of the RAs are valid by assumption, we only need to consider the validity of advertisements received from follower GAs at level $m + 1$. In particular, the (aggregated) belief set should represent correctly the uncertainty at the PCC during the entire computation cycle at m . However, this is ensured by construction described in Sections III-B and III-C. Indeed, a follower GA makes sure to compute setpoints that are consistent with the recently advertised belief set; see (2). Hence, the belief set is always valid, no matter what the value of Δ_t^m is. \square

IV. NUMERICAL ILLUSTRATION

In this section, we illustrate the performance of the proposed aggregation methods on a case study that makes reference to the low voltage microgrid benchmark defined by the CIGRÉ Task Force C6.04.02 [15], connected to a generic medium voltage feeder. This is the same case study used in [11] and is shown in Figure 1 (a). It is an islanded system, with a battery at the MV level (ESS resource in Figure 1 (a)) serving as a slack.

The simulation was performed using similar data and scenario as in [11]. In particular, nodal injections/absorptions of renewable resources are defined from experimentally inferred time series measured on the EPFL campus. The main difference from [11] is the more volatile profile for the LV uncontrollable load and solar radiation, which was created by adding artificial noise to the time series. This can be seen in Figure 8.

Figures 3-5 show the evolution of the aggregated PQ profile and belief set, during 4 seconds of simulation. It can be seen that the LVGA advertises a consistent advertisement and steers the LV grid according to the request of the MVGA. In turn, this request is consistent with the aggregated cost function shown in Figure 6. Indeed, the partial derivative with respect to P is positive, and hence the LVGA wants to minimize production (or maximize consumption) at this period. It can be seen in Figure 4 that the requested and actually implemented setpoints at the PCC (blue star and circle, respectively) are gradually decreasing as a result.

In Figures 7-9, the performance of the whole system is shown on a more larger time scale of 80 seconds. It can be seen that, as expected, the Commelec framework steers the state within the feasibility margins, minimizing the curtailment of the renewables and steering the storage devices towards their desirable values of SoC (in this case, set to 0.5). For more

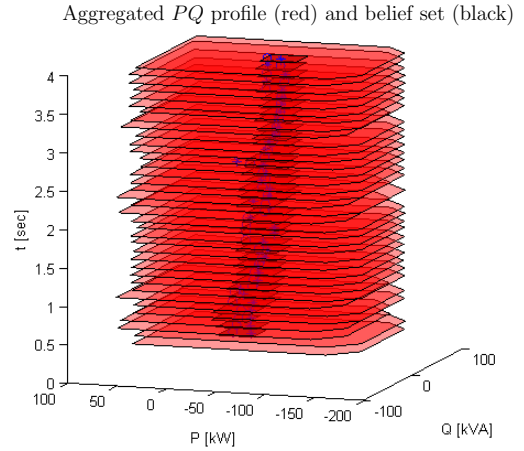


Figure 3. Aggregated PQ profile and belief set as a function of time computed by the LVGA. Blue star represents the requested power setpoint received from MVGA. Blue circle represents the implemented power setpoint at the PCC.

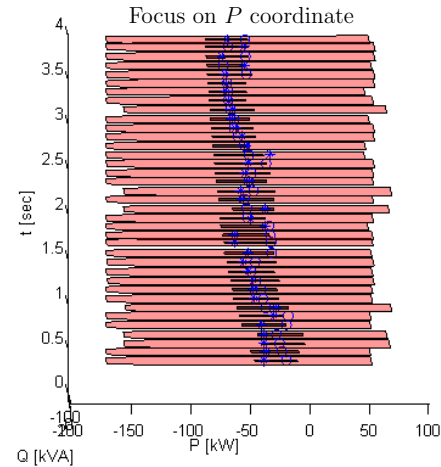


Figure 4. Aggregated PQ profile and belief set as a function of time computed by the LVGA – focus on the active power. Blue star represents the requested power setpoint received from MVGA. Blue circle represents the implemented power setpoint at the PCC.

details on the benefits of using Commelec in this case study see [11].

V. CONCLUSION

We proposed a computationally efficient method for aggregating heterogeneous electrical resources in the context of the Commelec framework for real-time control of medium and low voltage power grids. The method does not require synchronized operation, inherently robust to delays in communication, and always keeps the system in feasible electrical state. Another main characteristic of the proposed method is its applicability to systems with little or no inertia. Hence, the resulted aggregation provides higher flexibility than traditional approaches based on the assessment of aggregated inertia. We illustrated the performance of the method on a Cigre microgrid benchmark where realistic nodal power injections/absorptions

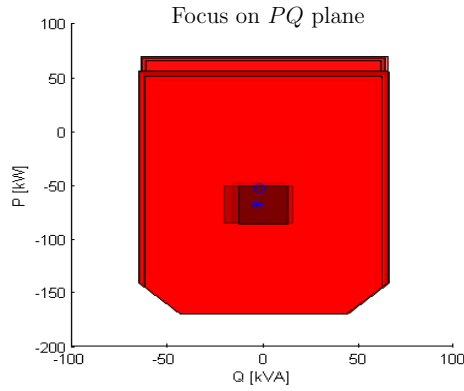


Figure 5. Aggregated PQ profile and belief set as a function of time computed by the LVGA – focus on the PQ plane. Blue star represents the requested power setpoint received from MVGA. Blue circle represents the implemented power setpoint at the PCC.

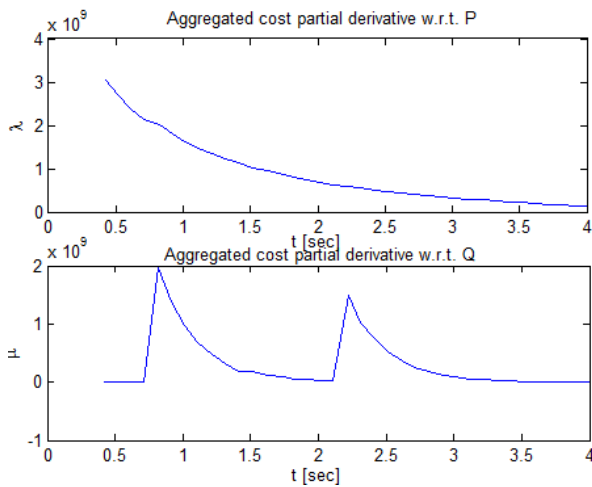


Figure 6. Derivatives of the aggregated cost function.

have been considered. To the best of our knowledge, this is the first method in the literature that allows a formal aggregation compatible with subsecond real-time applications.

REFERENCES

- [1] S. Tindemans, V. Trovato, and G. Strbac, "Decentralized control of thermostatic loads for flexible demand response," *Control Systems Technology, IEEE Transactions on*, vol. 23, no. 5, pp. 1685–1700, Sept 2015.
- [2] J. Mathieu, M. Kamgarpour, J. Lygeros, G. Andersson, and D. Callaway, "Arbitrating intraday wholesale energy market prices with aggregations of thermostatic loads," *Power Systems, IEEE Transactions on*, vol. 30, no. 2, pp. 763–772, March 2015.
- [3] S. Vandael, B. Claessens, M. Hommelberg, T. Holvoet, and G. Deconinck, "A scalable three-step approach for demand side management of plug-in hybrid vehicles," *Smart Grid, IEEE Transactions on*, vol. 4, no. 2, pp. 720–728, June 2013.
- [4] K. De Craemer, S. Vandael, B. Claessens, and G. Deconinck, "An event-driven dual coordination mechanism for demand side management of phev," *Smart Grid, IEEE Transactions on*, vol. 5, no. 2, pp. 751–760, March 2014.
- [5] B. Claessens, S. Vandael, F. Ruelens, K. De Craemer, and B. Beusen, "Peak shaving of a heterogeneous cluster of residential flexibility carriers using reinforcement learning," in *Innovative Smart Grid Technologies Europe (ISGT EUROPE), 2013 4th IEEE/PES*, Oct 2013, pp. 1–5.

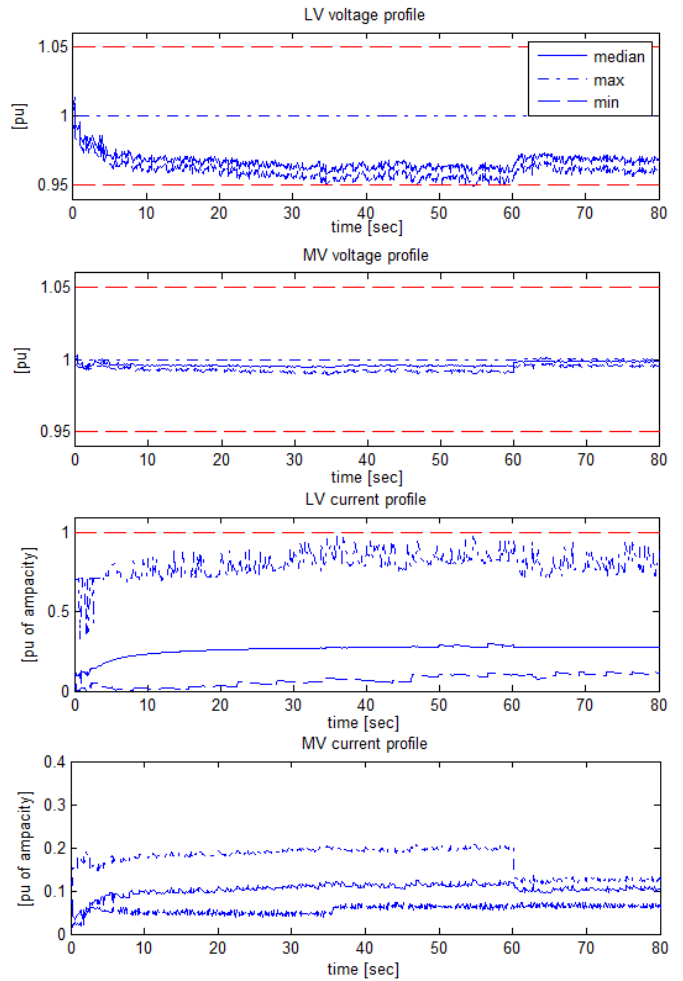


Figure 7. Voltage and current profiles.

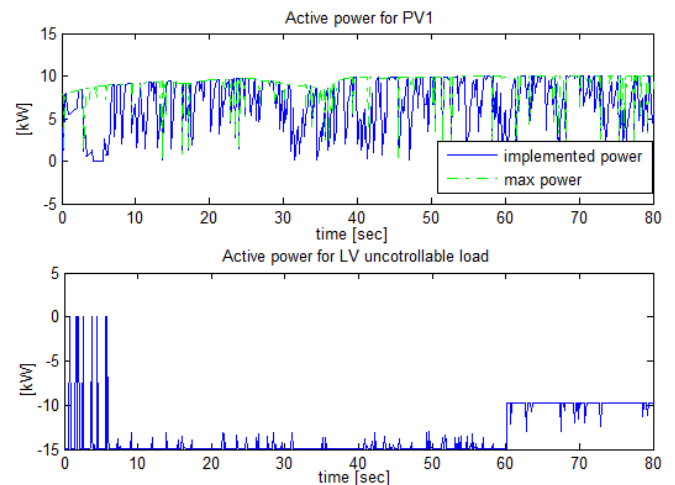


Figure 8. Active power of uncertain resources at the LV level: PV and uncontrollable load.




Figure 9. Power and energy of physical storage devices and demand response reservoirs.

- [6] F. Ruelens, B. Claessens, S. Vandael, B. De Schutter, R. Babuska, and R. Belmans, “Residential Demand Response Applications Using Batch Reinforcement Learning,” *ArXiv e-prints*, Apr. 2015.
- [7] M. Gonzalez Vaya and G. Andersson, “Optimal bidding strategy of a plug-in electric vehicle aggregator in day-ahead electricity markets under uncertainty,” *Power Systems, IEEE Transactions on*, vol. 30, no. 5, pp. 2375–2385, Sept 2015.
- [8] R. Pedersen, J. Schwensen, B. Biegel, J. Stoustrup, and T. Green, “Aggregation and control of supermarket refrigeration systems in a smart grid,” in *Proceedings of the 19th IFAC World Congress*. Cape Town, South Africa: IFAC, Aug. 2014. [Online]. Available: http://www.control.aau.dk/~jakob/selPubl/papers2014/ifacwc_2014_2.pdf
- [9] B. Biegel, P. Andersen, J. Stoustrup, M. B. Madsen, L. H. Hansen, and L. H. Rasmussen, “Aggregation and control of flexible consumers - a real life demonstration,” in *Proceedings of the 19th IFAC World Congress*. Cape Town, South Africa: IFAC, Aug. 2014. [Online]. Available: http://www.control.aau.dk/~jakob/selPubl/papers2014/ifacwc_2014_6.pdf
- [10] 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*, vol. 125, pp. 254–264, 2015.
- [11] 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*, vol. 125, pp. 265–280, 2015.
- [12] A. G. Expósito, A. Abur, A. de la Villa Jaén, and C. Gómez-Quiles, “A multilevel state estimation paradigm for smart grids,” in *Proceedings of the IEEE*, vol. 99, no. 6, pp. 952–976, June 2011.
- [13] M. Paolone, M. Pignati, P. Romano, S. Sarri, L. Zanni, and R. Cherkaoui, “A hardware-in-the-loop test platform for the real-time state estimation of active distribution networks using phasor measurement units,” *Proceedings of the CIGRE SC C6 Colloquium in Yokohama, Japan, October 6-9, 2013*.
- [14] 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 Transactions on Smart Grid*, vol. 4, no. 2, pp. 741–750, June 2013.
- [15] S. Papathanassiou, N. Hatziaegyriou, and K. Strunz, “A benchmark low voltage microgrid network,” in *Proceedings of the CIGRÉ Symposium*