Hierarchical Model Predictive Control for Resource Distribution

Jan Bendtsen Klaus Trangbaek Jakob Stoustrup

***Abstract*— This paper deals with hierarchical model predictive control (MPC) of distributed systems. A three-level hierarchical approach is proposed, consisting of a high level MPC controller, a second level of so-called *aggregators*, controlled by an online MPC-like algorithm, and a lower level of autonomous units. The approach is inspired by smart-grid electric power production and consumption systems, where the flexibility of a large number of power producing and/or consuming units can be exploited in a smart-grid solution. The objective is to accommodate load variations on the grid, arising from varying consumption and natural variations in power production, e.g. from wind turbines.**

The approach presented is based on quadratic optimisation and has low algorithmic complexity as well as good scalability. In particular, the proposed design methodology facilitates plug-and-play addition of subsystems without controller redesign. The method is verified by simulating a three-level smart-grid power control system for a small isolated power grid.

I. Introduction

We discuss a hierarchical setup, where an optimisation-based high-level controller is given the task of following a specific externally generated trajectory of consumption and/or production of a certain resource. The high-level controller has a number of units under its jurisdiction, which consume a certain amount of the resource. The flow of resources allocated to each of these units can be controlled, but each unit must at all times be given at least a certain amount of the resource; vice versa, each unit can only consume a certain (larger) amount of the resource.

One can think of various practical examples of systems that match this setup; for instance a supply chain management system (see e.g. [1]), where the challenge is to balance demand and supply using a number of storages with a maximal capacity each. The algorithm will then try to balance the risk of individual storages running empty or full with the risk of having over-production or unsatisfied demand. Other examples include large-scale refrigeration systems (e.g., in supermarkets), where the resource is refrigerant and the consuming units are individual display cases [2]; irrigation systems, where the shared resource is water and the consuming units are adjustable field sprinklers [3]; chemical processes requiring process steam from a common source [4]; or even digital wireless communication systems, where the resource is bandwidth and the consuming units are hand-held terminals, e.g. connected to a building-wide intranet [5].

Such large-scale hierarchical systems are often subject to frequent modifications in terms of subsystems that are added (or removed). This adds an important side constraint to design methodologies for controlling such systems: They should accommodate incremental growth of the hierarchical system in a way that is flexible and scalable. In essence, the design methodology should support a *plug-and-play control* architecture, see e.g. [6].

In this paper we propose a hierarchical control architecture that

- is based on a standard MPC solution at the top level
- remains stable for an increasing number of units
- facilitates plug-and-play of units at the bottom level, i.e., new units can be incorporated at the bottom level simply by registering the unit at the level just above it without requiring modifications of the top-level controller

We illustrate the approach by a specific example, a simple smart grid electric power system, where consumers can vary their power consumption within certain bounds by allowing devices to store more or less energy at convenient times, see e.g., [7].

The top level system should be flexible enough to accommodate new consumers under its jurisdiction without having to perform significant re-tuning and/or restructuring every time new consumers appear. Furthermore, it is a basic requirement that the system be stable and provide good performance at all times.

The outline of the rest of the chapter is as follows. Section II explains the problem in a general setting,

This work is supported by The Danish Research Council for Technology and Production Sciences.

The authors are with the Department of Electronic Systems, Automation and Control, Aalborg University, Denmark; email: {dimon, ktr, jakob}@es.aau.dk

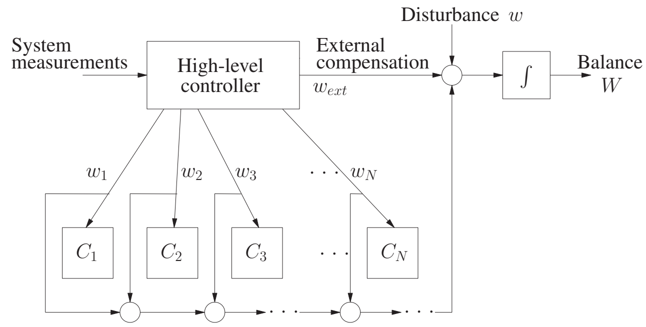

Fig. 1. Problem setup. The high-level controller must track a balance reference W while distributing resources to consumers C_i.

while Section III presents the proposed algorithm for resource sharing. Section IV shows that the resulting architecture remains stable for increasing numbers of units. Section V shows a simulation example of the algorithm applied to an electric smart grid with a small number of consumers, and finally Section VI offers some concluding remarks.

II. Problem Formulation

We consider a setup as depicted in Figure 1.

The high-level controller is given the task of following a specific externally generated trajectory of consumption and/or production of a certain resource. The objective is to maintain a certain system-level *balance* (between demand and production); the error in the balance is represented by the scalar signal[1] W, which must be driven to 0. Over time the demand and production must match, however, and the disturbance w is hence treated as short-time changes in the balance, whereas W is an integrated error signal. The high-level controller can compensate directly for the disturbance w by assigning some amount of the resource w_{ext} to this task, but at a significant cost. Furthermore, the high-level controller has a number of units C_i, which we will in general refer to as *intelligent consumers*, under its jurisdiction. Each of these consumers consume a certain amount of the resource w_i. The high-level controller is able to direct time-varying resources to the consumers, but must ensure that each consumer *on average* receives a specific amount of the resource, and certain upper and lower bounds $\underline{w}_i$ and $\overline{w}_i$ may not be exceeded. By doing so, the consumption compensates for some of the disturbance w, at a *lower cost* than the direct compensation signal w_{ext}. That is, it is advantageous to utilise the consumers as much as possible, subject to the aforementioned constraints.

[1]Note that, unless otherwise stated, all signals throughout the chapter are continuous-time and scalar.

In the following, let $\mathcal{I} = \{1, 2, \ldots, N\}$ denote an index set enumerating the consumers. The high-level controller must solve the following optimization problem at any given time t:

$$\begin{aligned} \min_{w_i, w_{\text{ext}}} \quad & \int_t^{t+N_h} \rho W(\tau)^2 + \phi(w_{\text{ext}}(\tau), \frac{dw_{\text{ext}}(\tau)}{dt}) d\tau \qquad (1) \\ \text{s.t.} \quad & \underline{W} \leq W(\tau) \leq \overline{W} \\ & \underline{w}_i \leq w_i(\tau) \leq \overline{w}_i, \quad i \in \mathcal{I} \end{aligned}$$

where $\underline{W}$ and $\overline{W}$ are constraints on the balance, ρ is a scalar cost on the balance deviation and $\phi : \mathbb{R} \times \mathbb{R} \to \mathbb{R}_+$ is a cost function of the absolute value of w_{ext} as well as changes in w_{ext}. N_h is the prediction horison of the controller. For simplicity, and without loss of generality, the consumption by the consumers is assumed cost-free.

Let W_i denote the amount of resource accumulated in C_i, and $\eta_i \geq 0$ denote a, possible time-varying, drain rate, respectively. Each consumer is characterised by its own linear state equation:

$$\frac{dW_i(t)}{dt} = w_i(t) - \eta_i(t) \qquad (2)$$

which must satisfy $0 \leq W_i(t) \leq \overline{W}_i$ at all times. Note that this model implies that the consumers are mutually independent. The goal that each consumer receives a specific amount of the resource on average, may be expressed as the integral constraint

$$\frac{1}{T_{\text{res}}} \int_0^{T_{\text{res}}} |w_i(\tau) - \eta_i(\tau)| d\tau = W_{i,\text{ref}} \qquad (3)$$

where T_{res} is some appropriate time span.

Note that, since the dynamics contain only pure integrators, (1) can easily be approximated by a discrete time problem

$$\begin{aligned} \min_{w_i, w_{\text{ext}}} \quad & \sum_{k=t/T_s+1}^{(t+N_h)/T_s} \rho W_k^2 + \phi(w_{\text{ext},k}, w_{\text{ext},k-1}) \qquad (4) \\ \text{s.t.} \quad & \underline{W} \leq W_k \leq \overline{W} \\ & \underline{w}_i \leq w_{i,k} \leq \overline{w}_i, \quad i \in \mathcal{I} \end{aligned}$$

where T_s is the sampling time.

In order to solve the optimisation problem, the high-level controller in principle requires access to all states in the system, including the internal states W_i. This may lead to a very heavy communication load on distributed systems. Furthermore, the computational complexity of the optimisation problem grows rapidly with

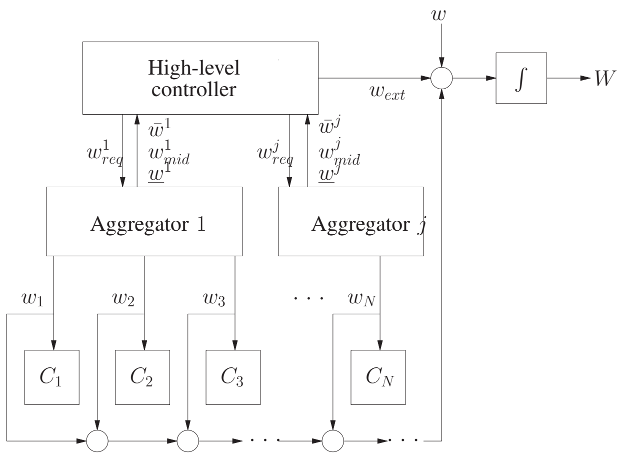

Fig. 2. Modified architecture

the number of consumers N. This means that adding more consumers into the system may pose significant problems in practice. Thus, a purely centralised solution to the problem may be optimal in terms of maintaining the supply/demand balance, but is not desirable from a practical point of view.

III. Proposed Algorithm

In the following we propose a new algorithm for achieving the control objective that requires significantly less system-wide communication, while at the same time being more flexible with respect to changes in the number of consumers.

We now consider the modified setup in Figure 2, where $w(t)$ is an external disturbance and $w_a(t) = \sum_{i=1}^N w_i(t)$ is the cumulative rate of resource absorbed by all C_i. As mentioned in the previous section, the main objective of the high-level control is to keep the resource balance governed by

$$\frac{dW(t)}{dt} = w(t) - w_{\text{ext}}(t) - w_a(t) \tag{5}$$

at zero. It is assumed that the top level controller can control $w_{\text{ext}}(t)$ directly and is constrained only by a rate limit, but we would like to keep the variations, i.e., the time derivative of $w_{\text{ext}}(t)$, small as well.

Between the controller and $N_A \leq N$ subsets of the intelligent consumers, we introduce a number of so-called *aggregators* $A_j, 1 \leq j \leq N_A$. Together, these aggregators serve as an interface between the top level and the intelligent consumers. To each aggregator A_j we assign a number of consumers identified by an index set $\mathcal{J}^j \subset \mathcal{I}$, where for all $k, j = 1, \ldots, N_A$ we have $\mathcal{J}^j \cap \mathcal{J}^k = \emptyset, k \neq j$, and $\cup_{j=1}^{N_A} \mathcal{J}^j = \mathcal{I}$. In short the objective of each aggregator is to make sure that:

- The maximum capacity is available for the upper level at any time instance
- The load for consumers is distributed roughly uniformly over the number of consumers
- The deviation from the nominal consumption is minimised for each consumer
- The rate and capacity constraints for each consumer are not violated

The communication between the high-level controller is indicated on Figure 2; A_j provides the top level with simple parameters to specify the constraints of the consumers. In particular, the top level is informed of $\overline{w}(t)$ and $\underline{w}(t)$, upper and lower limits on

$$w_a^j(t) = \sum_{i \in \mathcal{J}^j} w_i(t)$$

that can be guaranteed over the horison N_l. These limits depend on both resource storage rate and limitations among the individual consumers, and as such depend in a complicated fashion on the horison length. In addition to the limits, A_j provides w_{mid}^j, a mid-ranging signal that informs the high-level controller which total resource rate would be most helpful in bringing the intelligent consumers under its jurisdiction close to their reference resource levels $W_{i,\text{ref}}$. The aggregator level thus attempts to maintain $w_a(t) = w_{\text{req}}^j(t)$ while the high-level controller, in turn, needs only solve the optimisation problem

$$\begin{aligned} \min_{w_{\text{req}}^j, w_{\text{ext}}} \quad & \sum_{k=1}^{N_h} \rho W_k^2 + \phi(w_{\text{ext},k}, w_{\text{ext},k-1}) \\ \text{s.t.} \quad & \underline{W} \leq W_k \leq \overline{W} \\ & \underline{w}^j \leq w_{\text{req},k}^j \leq \overline{w}^j, \quad 1 \leq j \leq N_A \end{aligned} \tag{6}$$

which is of significantly lower dimension than (1) as long as N_A is significantly smaller than N. In periods where the load is relatively steady, the high-level controller can prioritise maintaining the balance, and thereby increasing the short term resource reserves for future load changes.

At each sample, the aggregator A_j solves the simple optimisation problem

$$\begin{aligned} \min_{w_i, i \in \mathcal{J}^j} \quad & \sum (W_i(t+T_s) - W_{i,\text{ref}})^2 \\ \text{s.t.} \quad & \sum_i w_i = w_{\text{req}} \\ & \underline{w}_i \leq w_i(t) \leq \overline{w}_i \\ & 0 \leq W_i(t+T_s) \leq \overline{W}_i \end{aligned} \tag{7}$$

with $W_i(t + T_s) = W_i(t) + T_s w_i$, where T_s is the sampling time.

IV. Stability Analysis

For every $j = 1, \dots, N_A$, it is clearly possible to gather W_i, $i \in \mathcal{J}^j$ in state vectors $\mathbf{x}^j$, $w_i - \eta_i(t)$ in vectors $\mathbf{u}^j$ and consider w^j_{req} as a reference input. Then the system to be controlled by the aggregator can then be described as the discrete-time system

$$\mathbf{x}^j(k+1) = \mathbf{A}^j \mathbf{x}^j(k) + \mathbf{B}^j \mathbf{u}^j(k) \tag{8}$$

where $\mathbf{x}^j$ and $\mathbf{u}^j$ are constrained to appropriate convex sets $\mathcal{X}_j$ and $\mathcal{U}_j$. Similarly, for the high-level controller, we can write the interaction with the aggregator level as the discrete-time system

$$\hat{\mathbf{x}}(k+1) = \hat{\mathbf{A}}\hat{\mathbf{x}}(k) + \hat{\mathbf{B}}\hat{\mathbf{u}}(k) + \hat{\mathbf{G}}\mathbf{w}(k) \tag{9}$$

where $\hat{\mathbf{x}}(k)$ contains samples of $w^j_{\text{mid}}, j = 1, \dots, N_A$, $\hat{\mathbf{u}}(k)$ contains samples of w^j_{req}, and $\mathbf{w}(k)$ represents the external disturbance. $\hat{\mathbf{x}}$ and $\hat{\mathbf{u}}$ are constrained to appropriate convex sets $\mathcal{X}$ and $\mathcal{U}$ determined by the signals $\underline{w}^j \le w^j_{\text{req},k} \le \overline{w}^j, \quad 0 \le j \le N_A$.

In the above, $\mathbf{A}^j, \mathbf{B}^j, \hat{\mathbf{A}}, \hat{\mathbf{B}}$ and $\hat{\mathbf{G}}$ are appropriately defined constant matrices. Note also that there is no difficulty in assuming that the high-level controller operates at a lower sampling rate than the aggregator level; in that case, we simply require the control signal $\mathbf{u}\hat{(k)}$ to be constant over the slow sampling period, i.e.

$$\hat{\mathbf{u}}(h\nu + \mu) = \hat{\mathbf{u}}(h\nu), \quad \mu = 0, \dots, \nu - 1$$

where h is the sampling rate of the high-level controller and $k = h\nu$ is the sampling rate of the aggregator level.

Formulated in this manner, the hierarchical control system can be seen to be exactly on the form considered in [8]; thus, we can invoke the main results in that paper and conclude the following (assuming the original problem is feasible):

1) Since all the individual subsystems are open loop (marginally) stable, all trajectories will tend to constant values for constant inputs.
2) In steady state, the minimal number of constraints are active. This follows from properties of quadratic optimisation; if the number of active constraints is non-minimal, the quadratic cost will always become smaller by shifting load from one of the subsystems with an active constraint to one or more subsystems with inactive constraints.
3) Since all the individual subsystems are open loop (marginally) stable, wind-up behaviour can be avoided for any bounded input set.

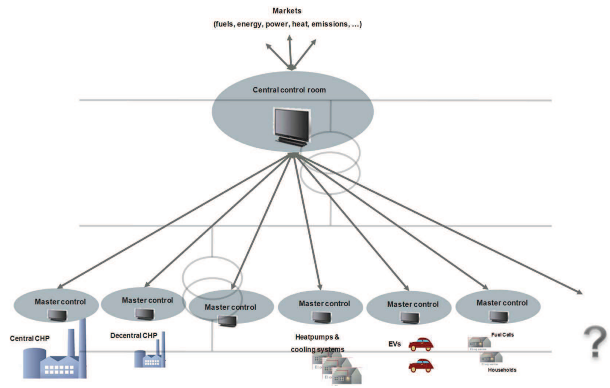

Fig. 3. A vision for Smart Grids: Virtual Power Plants which aggregate producing or consuming units.

V. Simulation Example

The example described in this section is inspired by a vision for future Smart Grid technologies called Virtual Power Plants, which is depicted in Figure 3.

The main objective of the top level control is to keep the energy balance governed by

$$\frac{dE(t)}{dt} = P_{\text{ext}}(t) - P_{\text{load}}(t) - P_{\text{a}}(t) \tag{10}$$

at zero. $P_a = \sum_i P_i$ is the power absorbed by the intelligent consumers (ICs). P_{load} is the power absorbed by other consumers, and is considered as a disturbance here. P_{ext} is the power produced by a number of suppliers such as power plants etc. It is assumed that the top level controller can control P_{ext} directly and is only restrained by a rate limit, but we would also like to keep the time derivative small.

Each IC is characterised by its own energy balance

$$\frac{dE_i(t)}{dt} = P_i(t) \tag{11}$$

which must satisfy $0 \le E_i(t) \le \overline{E}_i$ at all times. Furthermore, each IC is only able to consume a limited amount of power $\underline{P}_i \le P_i(t) \le \overline{P}_i$.

A set of aggregators serves as an interface between the top level and the ICs. It attempts to maintain $P_a(t) = P_{\text{req}}(t)$ and provides the top level with simple parameters to specify the constraints of the ICs. In particular, the top level is informed of $\overline{P}$ and $\underline{P}$, upper and lower limits on P_a that can be guaranteed over the horison N_l. In addition to the limits, the aggregators provide P_{mid}, a mid-ranging signal that tells the top level which P_{req} would be most helpful in bringing the ICs close to their reference energy levels $E_{ref,i}$.

How to choose these reference levels is again a complicated question of the considered horison. If we consider a long horison, then we might like to have the same energy reserve in both directions, which would lead to $E_{ref,i} = \overline{E}_i/2$. On the other hand, some ICs have a much higher $\overline{P}$ than $-\underline{P}$, and are therefore much better at providing a positive than negative absorption, while others are better at providing negative absorption. On a short horison it would make sense to keep the first kind at a low energy level, and vice versa. Here we choose $E_{ref,i} = \overline{E}_i \frac{\overline{P}_i}{\overline{P}_i - \underline{P}_i}$, which corresponds to making the time to fully charging identical to the time required to empty the energy reserve.

At each sample, at time t, the aggregator solves the simple optimisation problem

$$\begin{aligned} \min_{P_i} \quad & \sum \quad (E_i(t+T_s) - E_{i,\text{ref}})^2, \\ & \text{s.t.} \\ \sum P_i \quad &= \quad P_{\text{req}}, \\ \underline{P}_i \quad &\leq \quad P_i(t) \leq \overline{P}_i, \\ 0 \quad &\leq \quad E_i(t+T_s) \leq \overline{E}_i \end{aligned}$$

with $E_i(t+T_s) = E_i(t) + T_s P_i$, thereby distributing the power in a way that brings the energy levels as close to the reference as possible in a quadratic sense.

The top level control optimises over a prediction horison N_p. It minimises the performance function

$$\begin{aligned} J_t = & \sum_{k=1}^{N_p} E(t+T_s k)^2 \\ & + \beta_p \sum_{k=1}^{N_c} (P_{\text{ext}}(t+T_s k) - P_{\text{ext}}(t+T_s(k-1)))^2 \\ & + \beta_r \sum_{k=1}^{N_c} (P_{\text{req}}(t+T_s k) - P_{\text{mid}}(t))^2 \end{aligned}$$

with N_c samples of P_{ext} and P_{req} as decision variables.

The optimisation is subject to constraints on the decision variables. There is a rate limit on the power from the power plants:

$$\underline{P}_{\text{ext}} \leq P_{\text{ext}}(t+T_s k) - P_{\text{ext}}(t+T_s(k-1)) \leq \overline{P}_{\text{ext}}$$

As mentioned, the aggregator provides limits on P_a that can be sustained over a horison N_l. These limits are conservative in the sense that if P_{req} is for instance negative for the first part of the horison, then a positive P_{req} higher than $\overline{P}$ may be feasible for the rest. However, in order to simplify the top level computations, the constraint $\underline{P}(t) \leq P_{\text{req}}(t+T_s i) \leq \overline{P}(t)$ is imposed over the whole horison.

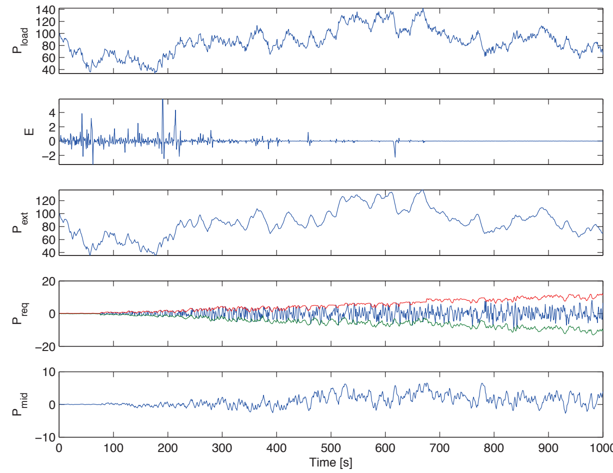

Fig. 4. Simulation example, top-level signals. As ICs are added (see Fig. 5), the control system's ability to compensate for external disturbances improves, as can be seen from the second plot. P_{mid} is computed as the P_{req} that would bring the energy levels to the reference in N_l samples, ignoring power limits.

A simulation of this scheme is shown in Figure 4. The controller parameters used are $T_s = 1$, $N_l = N_c = 4$, $N_p = 8$, $\beta_p = 0.1$, $\beta_r = 10^{-4}$. The load is generated by a first order auto-regressive process with a time constant of 100 seconds. There are 20 ICs with parameters shown in Table I becoming available as time passes, making it possible for the aggregator to provide wider constraints on P_{req}. The result is that the energy balance can be controlled much better while also using a smoother P_{ext}. The requested consumption P_{req} is shown together with $\underline{P}(t)$ and $\overline{P}(t)$, computed by the aggregator. It is noted how the constraints widen as more ICs become available, but will shrink when the reserve is being used.

The energy balance of the ICs is shown in Figure 5. The energy constraints and reference are shown by dashed lines. It can be seen that additional consumers are "plugged in", the system automatically incorporates these new consumers and these new resources are exploited throughout the control hierarchy in order to improve the power balance at the top level.

VI. Discussion

In this chapter a design methodology for a three level hierarchical control architecture is proposed. The emphasis is on systems that accumulate the production and/or consumption of resources through the levels,

i	$\overline{E}_i$	$\underline{P}_i$	$\overline{P}_i$	i	$\overline{E}_i$	$\underline{P}_i$	$\overline{P}_i$
1	1.0	-1.7	1.4	11	9.0	-0.2	1.1
2	4.0	-1.4	0.8	12	1.0	-1.0	1.2
3	4.0	-0.2	1.8	13	2.0	-1.6	1.6
4	3.0	-1.3	0.3	14	10.0	-1.3	1.9
5	6.0	-1.6	0.9	15	6.0	-0.3	0.4
6	10.0	-1.3	1.1	16	1.0	-1.1	0.9
7	1.0	-0.7	1.2	17	9.0	-1.9	1.8
8	4.0	-1.9	0.2	18	8.0	-0.2	0.8
9	9.0	-1.1	0.2	19	2.0	-0.9	0.6
10	10.0	-1.1	0.2	20	9.0	-1.6	0.5

TABLE I

PARAMETERS FOR 20 CONSUMERS IN SIMULATION

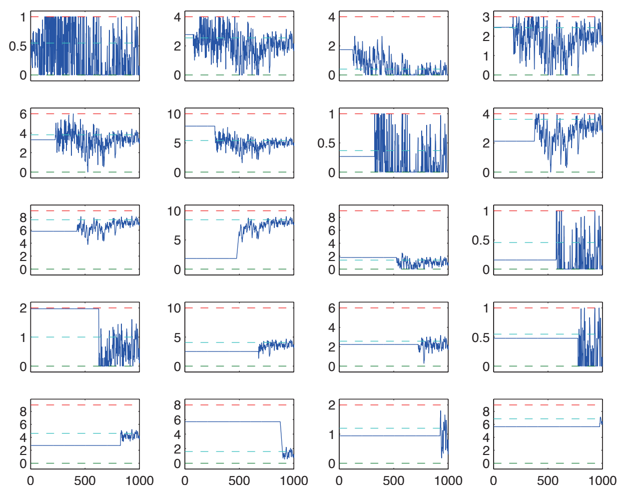

Fig. 5. Simulation with aggregators. Each plot shows the consumption of a specific IC. Each IC begins to provide consumption control capacity to the aggregator at varying time instances; prior to that point, their consumption rates are constant.

exemplified by irrigation systems, sewer systems, or power production and consumption systems.

The presented solution is based on MPC-like algorithms, based on online quadratic programming solvers. Considering the facts that the *total installed flexible capacity* (i.e., the sum of maximal resource storages for all units), *The instantaneous flexible capacity* (i.e., the currently unexploited part of C_{tot}), *The total cumulative rate limitation of flexible units* (i.e., the rate limitation experienced by the high level controller) and the *The instantaneous cumulative rate limitation of flexible units* (i.e., the current rate limitation experienced by the high level controller) all scale linearly with the number of ICs, the overall algorithmic complexity scales approximately with the number of units in the system to the power of 1.5, even without exploiting a significant sparsity of the optimization problems involved. Rigorous performance and complexity evalutions are subjects of future research, however.

The approach has the specific feature that it facilitates online modifications of the topography of the controlled system. In particular, units at the lower level can be added or removed without any retuning of any controllers. This plug-and-play control property is enabled by the modular structure of the involved cost functions of the optimisations.

The proposed methodology is exemplified by a simulation of a control system for a small electrical power production and consumption system, where the power flexibility of a number of consumers is exploited.

REFERENCES

[1] A. De Kok and S. Graves, “Supply chain management: Design, coordination and operation,” in *Handbooks in Operations Research and Management Science: Volume 11*, A. De Kok, Ed. Elsevier, 2003.

[2] L. Zheng, L. F. S. Larsen, R. Izadi-Zamanabadi, and R. Wisniewski, “Analysis of synchronization of supermarket refrigeration system- benchmark for hybrid system control,” *Kybernetes*, 2009.

[3] Y. Zhang, J. Zhang, and J. Guan, “Study on precision water-saving irrigation automatic control system by plant physiology,” in *Proc. of 4th IEEE Conference on Industrial Electronics and Applications*, 2009, pp. 1296–1300.

[4] K. Karlsson and I. Engstrom, “Controlling power and process steam,” *International Power Generation*, vol. 26, p. 29, 2003.

[5] C. J. Beckmann, “Modeling integrated services traffic for bandwidth management in next-generation intranets,” in *Proceedings of the 1997 5th International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems*, 1997.

[6] J. Stoustrup, “Plug & play control: Control technology towards new challenges,” *European Journal of Control*, vol. 15, no. 4, pp. 311–330, 2009.

[7] K. Moslehi and R. Kumar, “A reliability perspective of the smart grid,” *IEEE Transactions on Smart Grid*, vol. 1, no. 1, pp. 57–64, 2010.

[8] R. Scattolini and P. Colaneri, “Hierarchical model predictive control,” in *Proc. of the 46th IEEE Conference on Decision and Control*, Dec. 2007, pp. 4803–4808.