European Conference Data Mining 2014 and
International Conferences Intelligent Systems and Agents 2014 and
Theory and Practice in Modern Computing 2014

HOLISTIC OPTIMIZATION OF HVAC SYSTEMS VIA DISTRIBUTED DATA-DRIVEN CONTROL

Karel Macek[1], Dimitrios Rovas[2], Mischa Schmidt[3] and Cesar Cesar Valmaseda[4]
[1]*Honewell ACS Global Labs – Prague, V Parku 18, Prague 4, 148 00, Czech Republic, Institute of Information Theory and Automation, AVCR, Pod Vodárenskou věží 4, 14800, Prague 8, 18000, Czech Republic*
[2]*Technical University of Crete, University Campus, 73100 Chania, Greece, Fraunhofer Institute for Building Physics IBP, Gottschalkstrasse 28a, 341 27 Kassel, Germany*
[3]*NEC Laboratories Europe, Kurfürsten-Analage 36, 69115 Heidelberg, Germany*
[4]*Cartif, Parque Tecnológico de Boecillo, 205, 47151 Boecillo, Spain*

ABSTRACT

In the present paper, the control design of Heating, Ventilation and Air Conditioning (HVAC) systems is investigated. In large-scale buildings – e.g. hotels or hospitals – the high dimension of the control design problem precludes a solution with reasonable computational effort. In this paper, a distributed control strategy is proposed, where interacting agents are operating sub-systems; interaction between these agents can ensure that an optimum solution can be obtained. A novel method to distributed control tis introduced based on data-driven modeling where the strategy is not based on explicit optimization, but on weighted learning of the control rules; two examples of the addressed system are formulated. A significant advantage of the proposed approach consists in minimal assumptions on the addressed system and the most significant disadvantage is the need of sufficiently rich data-sets.

KEYWORDS

HVAC systems, distributed control, data-driven control, lazy learning, dynamic programming.

1. INTRODUCTION

The present paper focuses on the control of HVAC (heating, cooling, and air conditioning) systems. Recent research addresses the segment extensively, with motivation largely stemming from the fact that 54% of the overall consumption of energy in buildings is used for the operation of HVAC systems (Mathews et al. 2001). While the benefits of effective energy utilization can be sizable, the application of advanced control in this industry remains limited. We can observe several direction of the interest related to the energy savings in buildings: (i) good energy-management concepts and material selection, where the savings are achieved by the installation of latest devices and use of high-tech insulation materials (Fokaides and Papadopoulos 2014); (ii) the determination of critical faults and inefficiencies that lead to corrective actions in terms of hardware replacement and maintenance (Kukal et al. 2009; Berka and Macek 2011); (iii) the efficient control of particular subsystems, e.g. in the heating season, the indoor zone temperature is kept at the lower level of comfort (Zhou et al. 2014), and – finally - (vi) holistic optimization of the whole building where individual subsystems cooperate in order to minimize the overall costs for a given prediction horizon.

A very natural way to the holistic optimization is to formulate a Model Predictive Control (MPC) problem using the system of the whole building and to solve it consequently using either means of mathematical programming (Prívara et al. 2011) or using soft-computing methods (Kontes et al. 2012; Macek et al. 2013). The large-scale MPC problems can be formulated easily using the machine readable description of the building such as IFC BIM (Cerovsek 2011) which makes the solutions flexible and modular.

The drawback of the standard MPC approach lies in the high-dimensionality of the considered models and the difficulty in constructing these models. So while a large number of studies have been performed on the potential of MPC, little has been done in the real-world . The situation is more egregious for larger-scale buildings such as hospitals or hotels having sometimes hundreds of zones and tens of generation units, being

ISBN: 978-989-8704-10-8 © 2014

equipped by complex distribution system and influenced by various, possibly random factors; the MPC control design problem becomes impossible to solve with existing computational resources.

The practical implementation of the holistic building control can be achieved either by (i) distinguishing the low-level from the supervisory-level control or by (ii) breaking down the overall task of holistic control into smaller (sub-)control tasks and solving these in an orchestrated way.

- In case (i), the supervisory-level relies on the low-level that is responsible typically for achieving of given set-points using PID controllers and the supervisory control optimizes these set-points only (Kontes et al. 2012; Macek et al. 2013).
- Approach (ii) tackles the holistic optimization in a distributed way where individual subsystems are controlled by agents, each optimizing their operation independently and communicate intermediate results with their neighbor agents for an orchestrated decision finding.

The HVAC control uses both models based on the established laws of physics as well as the black-box and data-driven models. The first approach is typically much more accurate in terms of description of the hardware of the HVAC system and requires significant configuration labor. The second approach is capable to deal with random factors and is easier to be installed. There are two basic approaches to the data-driven control: the first one consists in identification of the model of the system and consequent optimization of the control inputs; the second one transforms existing data to control rules using weighted learning.

In this paper a novel algorithm is presented to address the problem of distributed data-driven control using weighted learning. Section 2 provides the algorithm itself, starting from the centralized version for single-step control, discussing the multi-step control, and – finally – describing a multi-agent multi-step version. Section 3 formulates some illustrative examples and in Section 4 some concluding remarks are provided.

2. DISTRIBUTED DATA-DRIVEN CONTROL

2.1 Basic Notation

Throughout the paper, we will use $|i|$ as a number of elements; $p(\alpha|\beta)$ will denote a conditioned probability density function; $E(\alpha|\beta) = \int \alpha p(\alpha|\beta) d\alpha$ stands for the expected value; $N(\mu, \sigma^2)$ stands for (multivariate) normal distribution; **R** is the set of real numbers.

We are addressing the dynamic control with discretized time $t = 1,2, \ldots$ where the observable[1] states x_t and inputs $\mathrm{u_t}$ are real valued vectors and the single-step loss z_t is a real scalar. We assume that we have data $(u_\tau, x_\tau, z_\tau)_{\tau=1}^{t-1}$ and are about to decide about $u_t \in U$ so $x_t \in C$ satisfies the constraints C with some probability threshold P and the expected value $\mathrm{z_t}$ is possibly minimal. The system is assumed to behave in the probabilistic way:

$$x_t \sim p(x_t | x_{t-1}, u_t),$$

and

$$z_t \sim p(z_t | x_t, u_t).$$

Graphically, the model is described in Figure 1.

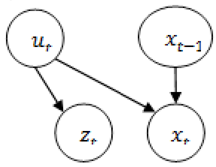

Figure 1. Single-step single-agent system.

[1] It is well known that the addressed HVAC systems contain important unobserved variables such as wall temperatures. Moreover, some states are observed with low precision, typically the outdoor air temperature. The assumption of fully observable state is adopted because of high novelty of the proposed solution.

Thus, the formulation of the control problem is

$$u_t^* = \arg\min_{u_t \in U} E[z_t | x_{t-1}, u_t]$$

subject to

$$\int_C p(x_t | x_{t-1}, u_t) dx_t \geq P.$$

2.2 Basic Algorithm

A data-driven approach to control was first presented by Atkeson et al. 1997 with the focus on reaching a given state x_t^{stp} rather than minimization of losses z_t. Two basic approaches are distinguished:

- the first consists in the creation of a data-driven model $x_t | x_{t-1}, u_t$ of the system and consequent explicit solution, e.g. based on gradient search so $x_t = x_t^{\mathrm{stp}}$. In the optimal control, this corresponds to the explicit minimization of the data-driven model of loss function $z_t | x_{t-1}, u_t$.
- the second approach is based on direct calculation of the inversion model $p(u_t | x_{t-1})$ based on a generalization of observed data where u_τ applied to $x_{\tau-1}$ led to x_τ^{stp}. In the optimal control, we have to determine the data where the control inputs applied to given state led to the minimal value of the loss function. In this case, the solution is more challenging since the selection of the relevant data has to reflect the fact that the minimal value of the loss function differ for different states $x_{\tau-1}$. Note that in the static optimization where no states are considered, the inversion is approximated by algorithms such as Covariance Matrix Adaptation (Hansen and Ostermeier 1996).

The idea of the algorithm is to generalize the past data using a – possibly simple – parameterized model for a decision rule $p(u_t | x_{t-1}, \theta)$. If we would assume that the data contain only records based on the explicit optimization, we can use usual Bayesian update, namely

$$p(\theta|\mathrm{data}) \propto p(\theta) \prod_{\tau=2}^{t-1} p(x_\tau | u_\tau, x_{\tau-1}) p(x_\tau | x_{\tau-1}, \theta) \propto p(\theta) \prod_{\tau=2}^{t-1} p(x_\tau | x_{\tau-1}, \theta);$$

where $p(\theta)$ is a prior distribution of the parameters θ. However, in practice, we have data that are not necessarily optimal: the actions were determined based on some suboptimal approaches or heuristics. The algorithm attempts to use the information in the data and to generalize it. Let us consider $w_\tau \in \mathbf{R}$ as a measure of the optimality of record τ. Then we can use

$$p(\theta|\mathrm{data}) \propto p(\theta) \prod_{\tau=2}^{t-1} [p(x_\tau | x_{\tau-1}, \theta)]^{w_\tau}$$

This can be calculated either online and in a batch and give us a decision rule that converges to the best approximation of the optimal decision rule with the given model structure. The only thing needed is to determine the weights w_τ as a function of $x_{\tau-1}$, u_τ, x_τ, and z_τ. First, let us introduce the extended loss function that involves also penalization of constraint violation, e.g.

$$\tilde{z}_\tau = z_\tau + \gamma[1 - \mathbf{1}(u_\tau \in U)] + \gamma[1 - \mathbf{1}(x_\tau \in C)]$$

where $\gamma > 0$ is a large number and $\mathbf{1}(A) = 1$ when the proposition A is true and $\mathbf{1}(A) = 0$ otherwise. The weights can be interpreted as the level of optimality with the system evolution in the given state $x_{\tau-1}$. Let us assume that we know the minimal and maximal value of the loss function that can be achieved for given $x_{\tau-1}$. Let us denote them $\tilde{z}_{\min}(x_{\tau-1})$ and $\tilde{z}_{\max}(x_{\tau-1})$ they can be obtained e.g. using Gaussian processes (Rasmussen and Williams 2005).

$$w_\tau = {}^I p(\tilde{z}_\tau | x_{\tau-1}) = \frac{\exp(-\alpha\, \tilde{z}_\tau)}{\int_{\tilde{z}_{\min}(x_{\tau-1})}^{\tilde{z}_{\max}(x_{\tau-1})} \exp(-\alpha \tilde{z}) d\tilde{z}}$$

Note that this is the normalized softmax (Sutton and Barto 1998). This formula shows the relatively lower loss function, the higher the corresponding weight. The normalization is motivated by the fact some records are very relevant for the learning of the control rule, but they started from very bad state $x_{\tau-1}$. One can use also some other efficiency measures, based e.g. on the DEA analysis (Banker et al. 1984). Also, more advanced approaches can be based on the transformation of the loss function to the ideal probability distribution (Kárný and Kroupa 2012).

Unfortunately, the Bayesian learning of probability densities of the parameters might seem relatively far away from the practical implementation. Let us therefore mention a well-established frequentist counterpart to the weighted Bayesian learning, namely the weighted least squares (Cleveland 1977) where the loss function

$$L(\theta) = \sum_{\tau=2}^{t-2} w_\tau \big(u_\tau - \hat{u}(x_{\tau-1},\theta)\big)^2$$

is minimized in θ where $\hat{u}(\cdot,\cdot)$ is a regression fit, e.g. the linear regression of $x_{\tau-1}$ with coefficients θ.

To summarize the algorithm, let us mention the following steps:

Offline/recursive

1. Calculate the weights w_τ.
2. Calculate the the control rule $p(\theta|\text{data})$ using the available data and the calculated weights w_τ.

Online

3. Apply the control rule $p(u_t|x_{t-1})$.

2.3 Extension to Multi-Step Control

The previous paragraphs have described a way how the data can be used for the calculation of the single-step control rule. Now, we extend the approach to the multi-step case. The basic idea is to start with the control rule $p(u_{t+|h|}|x_{t+|h|-1})$ for the final control rule of the decision horizon h, i.e. for time $t+|h|$.Then calculate $p(u_{t+|h|-1}|x_{t+|h|-2})$, then $p\big(u_{t+|h|-2}\big|x_{t+|h|-3}\big)$ and so on until $p(u_t|x_{t-1})$. The first rule $p(u_{t+|h|}|x_{t+|h|-1})$ is calculated as the single-step control rule as described in Section 2.2. The other rules are calculated using the loss-to-go information z_τ^{prev} from the previous rule. This is the value of being in x_τ when the h^{th} control input has to be determined.

1. Set $z_\tau^{\text{prev}} = 0$ for all $\tau = 1,2,\ldots t-1$ and set $h = |h|$.
2. Set $z_\tau^{\text{tmp}} = \tilde{z}_\tau + z_\tau^{\text{prev}}$ and calculate the weights using w_τ.
3. Using the weights w, calculate the decision rule $p(u_{t+h}|x_{t+h-1})$
4. Set $h = h-1$. If $h = 0$, go to 6., otherwise go to 5.
5. Calculate $z_\tau^{\text{prev}} = E\big(z_\tau^{\text{tmp}}\big|x_{\tau-1}\big)$, given the control rule $p(u_{t+h}|x_{t+h-1})$, using a regression model for $z_{\tau+1}^{\text{tmp}}|x_\tau, u_{\tau+1}$. Continue to 2.
6. Return the control strategy $\big(u_{t+|h|}\big|x_{t+|h|-1}\big),\ldots,p(u_t|x_{t-1})$.

2.4 Distributed Data-Driven Control

For distributing the data-driven control strategy we described in previous sections to sub-problems, we assume that the control actions can be decomposed into several groups; each shall be associated to an agent $u_t = \big(u_{t,1}, u_{t,2}\ldots, u_{t,|i|}\big)$. The decomposition with respect to the agents is described as follows:

- Loss function $z_t = \sum_{i=1}^{|i|} z_{t,i}$ can be decomposed as sum of particular subsystems.
- The components of the states are classified into 3 groups for each agent (Šmídl and Přikryl 2006):
  - States modeled by the agent $x_{t,\text{out}[i]}$,
  - states consumed by the agent $x_{t,\text{in}[i]}$,
  - and states neglected by the agent $x_{t,\text{no}[i]}$.

Whenever an agent has an input state, another agent has to have the same output state. The latter is denoted as influencer of the former agent. Each agent is assumed to have local history containing records $\big(x_{\tau,\text{in[i]}}, x_{\tau,\text{out[i]}}, u_{\tau,i}, \tilde{z}_{\tau,i}\big)_{\tau=1}^{t-1}$.

We assume the following dynamics:

$$x_{t,\text{out}[i]} \sim p(x_{t,\text{out}[i]}|x_{t-1,\text{in}[i]}, x_{t-1,\text{out}[i]}, u_{t,i})$$

The cost loss function is as

$$z_{t,i} \sim p(z_{t,i}|x_{t,\text{out}[i]}, u_{t,i})$$

Local decision rules in form $p(u_{t,i}|x_{t-1,\text{out[i]}}, x_{t-1,\text{in[i]}}, \theta_i)$ with the learning scheme

$$p(\theta_i|\text{data}_i) \propto p(\theta_i)\prod_{\tau=2}^{t-1}[p(x_\tau|x_{\tau-1},\theta_i)]^{w_\tau}$$

The algorithm for the multi-agent settings can be summarized in the following steps:

1. Set $z_{\tau,i}^{\text{prev}} = 0$ for all $\tau = 1,2,\dots t-1$ and for all $i = 1,2,\dots,|i|$ and set $h = |h|$.
2. Set $z_{\tau,i}^{\text{tmp}} = \tilde{z}_{\tau,i} + z_{\tau,i}^{\text{prev}}$ and calculate the weights $w_{\tau,i}$ using $z_{\tau,i}^{\text{tmp}}$ and α
3. Using the weights w, calculate the decision rule $p(u_{t+h,i}|x_{t+h-1,\text{out}[i]}, x_{t+h-1,\text{in}[i]})$
4. If $h = 0$, go to 8., otherwise set $h = h - 1$ and go to 5.
5. Provide feedback to all influencers $E(z_{t+h,i}|x_{t+h-1,j})$ where j is influencer of i.
6. Provide feedback to itself $E(z_{t+h,i}|x_{t+h-1,i})$.
7. Aggregate the feedback provided $E(z_{t+h}|x_{t+h-1}) = E(z_{t+h,i}|x_{t+h-1,i}) + \sum E(z_{t+h,j}|x_{t+h-1,i})$ where j are agents influenced by i.
8. and apply it to all records $z_{\tau,i}^{\text{prev}} = E(z_{t+h}|x_{t+h-1} = x_\tau)$
7. Calculate $z_\tau^{\text{prev}} = E\left(z_\tau^{\text{tmp}}|x_{\tau-1}\right)$, given the control rule $p(u_{t+h}|x_{t+h-1})$, using a regression model for $z_{\tau+1}^{\text{tmp}}|x_\tau, u_{\tau+1}$. Continue to 2.
8. Return the control strategies $\left(u_{t+|h|,i}|x_{t+|h|-1,i}\right),\dots,p(u_t|x_{t-1,i})$ for all $= 1,2,\dots,|i|$.

The algorithm is illustrated in Figure 2, including the feedback being provided to agent 2 at the time instant $t + |h| - 1$, i.e. $E(z_{t+|h|,1}|x_{t+|h|-1,1})$ and $E(z_{t+|h|,2}|x_{t+|h|-1,1})$.

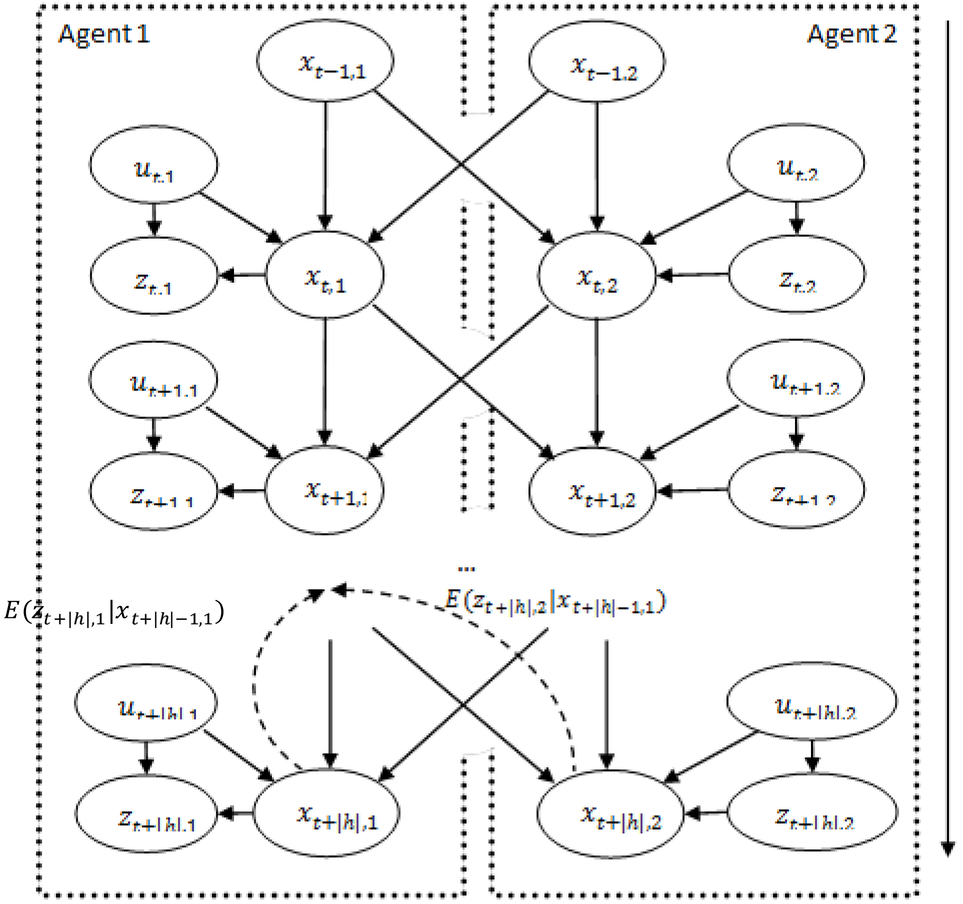

Figure 2. Example of multi-step control of two agents.

2.5 Holistic Optimization of Complex HVAC Systems

The distributed data-driven control proposed in the previous sections is a general framework that is applicable to wide class of building systems. Due to its flexible formulation in terms of cost functions, parameter spaces and the explicit interaction between agents, it is able to control heterogeneous subsystems reflecting their interaction while avoiding the need of formulating and solving a complex centralized optimization problem.

We can consider the following types of agents in the HVAC Systems are natural to express in our proposed approach:

- Generation/production – agents representing devices such as boilers, chillers, solar panels, co-generation units.
- Distribution – fans, dampers, pumps, valves, ducts.
- Demand – individual zones, including the individual preference for the comfort of the occupants.
- Others – weather forecasting agents.

The proposed algorithm is that the generation will be triggered by high penalty for comfort violation in the zones.

3. TOWARDS THE HOLISTIC HVAC OPTIMIZATION

In order to justify the applicability of the approach at least qualitatively, we formulate a simple HVAC related toy example in detail: the Optimal Heating in Two Zones. This example addresses the balancing of heating in two zones.

The states $\mathrm{x_t} \in \mathrm{R}^3$of the system involve two air temperatures in the zones and the outdoor air temperature. The control inputs $\mathrm{u_t} \in \mathrm{R}^2$ are the intensities of heating or cooling in each zone. We assume the dynamics of the system are

$$x_t \sim N(Ax_{t-1} + Bu_t, \Sigma_x)$$

With

$$A = \begin{bmatrix} 0.8 & 0.1 & 0.1 \\ 0.1 & 0.8 & 0.1 \\ 0 & 0 & 1 \end{bmatrix}$$

and

$$B = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix}.$$

Let us consider the heating in two zones where each zone is represented by one agent. We assume the extended loss function is for the first agent:

$$z_{t,1} = N(u_{t,1} + 100 \cdot \mathrm{dist}(x_{t,1}, C_1), \sigma_{z1})$$

while for the second agent

$$z_{t,2} = N(20 \cdot u_{t,2} + 100 \cdot \mathrm{dist}(x_{t,2}, C_2), \sigma_{z2})$$

where $\mathrm{dist}(\mathrm{s}, \mathrm{S})$ calculates the distance between vector s and set S and the comfort is given by intervals, i.e. $\mathrm{C_1} = \mathrm{C_2} = [20, 22]$. Thus, the heating in zone 2 is twice expensive than in the zone 1. We will consider the prediction horizon $|\mathrm{h}| = 10$. The initial state will be $\mathrm{x_0} = (21, 21,16)'$.

While yhis problem can be solved easily using centralized MPC based on linear programming (Garcia 1989), eventually by adopting the distributed MPC (Moroşan et al. 2011), it is possible to solve this toy example problem as well using the data-driven approaches based on inversion learning in both centralized and distributed settings as described in the present paper.

4. CONCLUSION AND FUTURE WORK

The proposed distributed data-driven control is a general and flexible framework that is applicable to wide class of systems and is able to approximate the optimal control of individual subsystems including their interaction without the need of centralized problem formulation and solution. The learned control strategy approaches the optimal strategy and does not require – in contrast to usual application of the MPC – iterative re-calculations. The agents are assumed to provide some loss-to-go feedback to their influencers for each time step of the prediction horizon.

The data-driven control based on weighted learning of control rules does not require description of the system in terms of first principles. The only information needed is the structure of influence among the agents. On the other hand, available first principle knowledge can be entered into the proposed framework as a part of the prior distributions of the sub-system parameters.

The basic principle is that the agents communicate the loss-to-go feedback to their influencers. This general principle can be approached either via the data-driven control based on weighted learning of control rules as described above or using any other approach. This makes it possible to integrate heterogeneous agents, including some special agents representing the forecasters of random effects.

The practical implementation of the algorithm for the HVAC control can be considered in various alternatives. First, very natural implementation would be when each agent[2] owns a dedicated processor and communicates e.g. via wireless network. However, since the calculation of the universal control strategy might be computationally intensive, it is also possible to carry out the calculations in the cloud or in a big-data infrastructure that is connected to the basic control infrastructure.

The proposed algorithm is not applicable only to HVAC system. The multi-agent settings motivates the integration also with related systems such a load dispatch, optimal maintenance, or water management. A further field of applicability of our proposed approach is to formulate the scheduling of the optimal maintenance actions for devices in time (Berka and Macek 2011). It has been shown that the problem can be formulated and solved in terms of dynamic programming. To illustrate the suitability of the distributed approach presented in this paper, let us consider two devices. The states are whether the devices are healthy or faulty $\mathrm{x_t} \in \{0,1\}^2$. The control inputs are the maintenance actions $\mathrm{u_t} \in \{0,1\}^2$. The dynamics can be described as a Markov chain where:

- $p\left(x_{t,i} = 0 \middle| u_{t,i} = 1\right) = 1$ i.e. after maintenance, the system becomes healthy.
- $p(x_{t,i} = 1 | u_{t,i} = 0, x_{t-1,i} = 1) = 1$, i.e. without maintenance, the system will not become healthy if it was faulty.
- We assume for this the risk of fault occurrence
    - $p(x_{t,1} = 1 | u_{t,1} = 0, x_{t-1,1} = 0, x_{t-1,2} = 0) = 0.02$
    - $p\left(x_{t,1} = 1 \middle| u_{t,1} = 0, x_{t-1,1} = 0, x_{t-1,2} = 1\right) = 0.10$
    - $p\left(x_{t,2} = 1 \middle| u_{t,2} = 0, x_{t-1,2} = 0, x_{t-1,1} = 0\right) = 0.01$
    - $p(x_{t,2} = 1 | u_{t,2} = 0, x_{t-1,2} = 0, x_{t-1,1} = 0) = 0.30$

We also assume the loss function is $z_t = 5u_{t,1} + 10u_{t,1} + 5x_{t,1} + 100x_{t,1}$. The values can be interpreted in the following way: the agent for device 1 neglects the faulty state or the maintenance. However, if considering the second device, it shall be rather healthy because it significantly reduces the risk of occurrence of very expensive fault at the second device. Note that the problem of the optimal maintenance can be combined with the optimal control.

Concerning challenges, the most serious open question from a research perspective is the practical testing of the proposed algorithm and its further theoretical justification. This task is connected to setting of the algorithm's parameters: including the selection of the $p(u_t | x_{t-1}, \theta)$ structure, to ensure the weights w_τ express sufficiently a level of optimality and whether the data-set is rich enough. Another task is a natural decomposition of the system, selection of the sampling period and definition of local loss functions.

A very interesting research challenge consists also in the determination to which extent and in which settings a centralized approach to optimization outperforms the proposed distributed approach and vice versa: The distributed approach is able to decompose very complex systems to computationally feasible problems. On the other hand, its intensive communication makes the approach inappropriate for simple systems. This is related also to the problem of granularity of the decomposition – i.e. whether it is better to use decomposition into many small subsystems with simple local optimization or whether to use the decomposition into several larger subsystems with more complex local optimization.

[2] Each agent can be represented by an APO (Assessment, Prediction, Optimization) unit (Valmaseda 2013).

ISBN: 978-989-8704-10-8 © 2014

It remains for future work to verify and assess the proposed approach both experimentally and in simulations in a variety of real-world settings.

ACKNOWLEDGEMENT

The research leading to this work has been supported by the European Commission FP7-2011-6, ICT Systems for Energy Efficiency under contract #288409 (BaaS – Building as a Service) and Grant Agency of the Czech Republic GACR 13-13502S.

REFERENCES

Atkeson, C. G. et al, 1997. Locally weighted learning for control. In *Lazy learning* (pp. 75-113). Springer Netherlands.

Banker, R. D. et al, 1984. Some models for estimating technical and scale inefficiencies in data envelopment analysis. *Management science*, Vol. 30, No. 9, pp. 1078-1092.

Berka J. and Macek K., 2011. Effective Maintenance of Stochastic Systems via Dynamic Programming. *Proceedings Technical Computing*. Prague, Czech Republic.

Cerovsek, T., 2011. A review and outlook for a 'Building Information Model' (BIM): A multi-standpoint framework for technological development. *Advanced engineering informatics*, Vol. 25, No. 2, pp. 224-244.

Fokaides, P. A., and Papadopoulos, A. M., 2014. Cost-optimal insulation thickness in dry and mesothermal climates: Existing models and their improvement. *Energy and Buildings*, Vol. 68, pp 203-212.

Garcia, C. E. et al, 1989. Model predictive control: theory and practice—a survey. *Automatica*, Vol. 25, No. 3, pp. 335-348.

Hansen, N., and Ostermeier, A., 1996. Adapting arbitrary normal mutation distributions in evolution strategies: The covariance matrix adaptation. *Evolutionary Computation*, Nagoya, Japan, pp. 312-317.

Kontes, G. D. et al, 2012. Adaptive-fine tuning of building energy management systems using co-simulation., *Control Applications*, Dubrovník, Croatia , pp. 1664-1669.

Kárný, M., and Kroupa, T., 2012. Axiomatisation of fully probabilistic design. *Information Sciences*, Vol. 186, No. 1, pp. 105-113.

Kukal, J. et al, 2009. From Symptoms to Faults: Temporal Reasoning Methods. *Adaptive and Intelligent Systems*. Klagenfurt, Austria, pp.155-159.

Macek, K. et al, 2013. Black-Box Optimization for Buildings and Its Enhancement by Advanced Communication Infrastructure. *ADCAIJ: Advances in Distributed Computing and Artificial Intelligence Journual*, Vol. 1, No. 5, pp 53-64.

Mathews, E. H. et al, 2001. HVAC control strategies to enhance comfort and minimise energy usage. *Energy and Buildings*, Vol. 33, No. 8, pp. 853-863.

Moroşan, P. D. et al, 2011. A distributed MPC strategy based on Benders' decomposition applied to multi-source multi-zone temperature regulation. *Journal of Process Control*, Vol. 21, No. 5, pp. 729-737.

Prívara, S. et al, 2011. Model predictive control of a building heating system: The first experience. *Energy and Buildings*, Vol. 43, No. 2, pp. 564-572.

Rasmussen, C. E. and Williams, C. K. I., 2005, *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. MIT Press, Cambridge, USA.

Šmídl, V., and Přikryl, J., 2006. Distributed Bayesian Decision-Making for Urban Traffic Control. *IEEE Industrial Electronics*, Paris, France, pp. 4695-4700.

Sutton, R. and Barto, A., 1998, *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, USA.

Valmaseda, C. et al, 2013. *An Event-driven SOA-based Platform for Energy-efficiency Applications in Buildings*, Applications of IT in AEC Industry, Beijing, China.

Zhou, P. et al, 2014. Demand-based temperature control of large-scale rooms aided by wireless sensor network: Energy saving potential analysis. *Energy and Buildings*, Vol. 68, pp. 532-540.