Case-based Reasoning for Inert Systems in Building Energy Management

Mirjam Minor and Lutz Marx

Goethe University, Business Information Systems, Robert-Mayer-Str. 10,
60629 Frankfurt, Germany
minor@cs.uni-frankfurt.de,
wirtschaftsinformatik@lmarx.de

Abstract. Energy management systems are a typical example for inert systems where an event or action causes an effect with a delay. Traditional solutions for energy management, such as PID controllers (PID = proportional-integral-derivative loops), control target values efficiently but are sub-optimal in terms of energy consumption. The paper presents a novel, case-based reasoning approach for inert energy management systems that aims to reduce energy wastage in over heating and over cooling for buildings. We develop a case representation based on time series data, taking environmental impact factors into consideration, such as weather forecast data. This includes a post-mortem assessment function that balances energy consumption with comfort for the users. We briefly discuss retrieval and reuse issues. We report on an experimental evaluation of the approach based on a building simulation, including 35 years of historical weather data.

Keywords: Case-based reasoning, reasoning over time, energy management

1 Introduction

A system can change its state by the influence of an impact factor. In physics or in biology, we necessarily have a time delay between impact and state change. *Inert Systems* are systems where an event or action takes effect with a delay and over a period of time. The system may either abruptly switch from one discrete state to another after a delay time or cumulate the impact factor and change the state continuously. In the following, we will focus on a case-based approach for the latter. An example of an inert system in nature is the human body where the injection of a drug changes the insulin level for a couple of hours. A technical sample is an energy management system (EMS) for buildings. The movement of a weather front has an impact on the room temperature with a time lag. We have chosen energy management for buildings as a sample application area for controlling inert systems.

Traditionally, Control System Engineering (CSE) [8] is used for this class of problems. It is well understood and widely used in EMSs and other domains.

PID controllers (proportional-integral-derivative loops) and switching rules are the industry standard for the control of building EMSs [6]. This logic responds to setpoints and schedules for building components, such as heating circuits, radiators, or air handling units. That means that the temporal delays between causes and effects that are characteristic for inert systems are only considered in a reactive manner by the controller. The temporal dependencies are hidden in setpoints and schedules. For instance, the time for pre-heating to change a heating circuit into an 'enabled' mode is expressed by higher setpoints in the early morning schedule. The basic control logic largely ignores forward planning based on weather forecasts, expected occupancy, or renewable energy availability.

More advanced, model-based decision systems aim to optimize the system operation based on modeling, feedback, and forecasts [6]. They use an explicit time model and forward planning in order to optimize the energy consumption at a system level. However, solving an optimization problem has two challenges [6] in comparison to the basic control logic, such as PIDs: It requires analytical building models at design time and, second, it is computationally intensive, i.e. it requires powerful computational units. A novel EMS is desirable with a lower energy consumption than a PID controller but that is easy to deliver, and easy to operate.

As an alternative solution to costly optimization, *Case-based reasoning* (CBR) provides methods for experience reuse. In this paper, we propose a CBR approach for energy management in buildings where experience in operating the EMS is to be reused. The traditional PID controllers of the EMS are replaced by a case-based control unit for the energy supply. Since EMSs for buildings are inert systems, the cases need to be equipped with a concept of time. The core idea is to observe and record the inert behavior of a system by time series of impact factors and state variables. The context description for a case comprises further time series, such as measured values of the building, recent metereological data and weather forecasts. Corrective actions that have led to good results in similar situations in the past are reused to manipulate the system state in the next time step, i.e. to achieve a system state that is close to the setpoint values. For instance, if a setpoint in an EMS specifies a desired room temperature a corrective action is an amount of energy to be supplied to or dissipated from the room. Like a PID controller that provides a corrective action as an output of each control cycle in order to maintain a desired setpoint value, our CBR approach provides a corrective action as an output of each reasoning cycle.

In comparison to the optimization approaches, the case-based approach uses a shallower model. The analytical model is built on similarity functions for cases. We claim that the CBR approach outperforms the basic control approach for building EMSs in terms of energy consumption while providing the same comfort for the occupants.

The remainder of the paper is organized as follows. Related work is discussed in Section 2. The case representation is introduced in Section 3. A similarity function is presented in Section 4. An adaptation rule is specified in Section 5.

Section 6 addresses the experimental setup while the results are reported in Section 7. Finally, a conclusion is drawn in Section 9.

2 Related work

CBR has been used for energy prediction in the recent literature [9, 11]. Like our approach, the work uses a notion of time series. The case-based energy prediciton approaches forecast the energy consumption based on energy values from cases with time series data that is similar to the recent situation. The prediction task is different but quite related to the control task that we address in our work. Like in prediciton, we reuse cases with similar time series. In contrast to reusing energy values, we reuse the control actions that have an impact to the inert system and observe the resulting energy consumption.

Temporal context plays a major role in many CBR approaches [10, 5]. Recent work on CBR on time series data is reported in Gundersen's survey [2] as well as in the series of RATIC workshops [4, 3]. This work was a major source of inspiration for our case representation and recent similarity function.

3 Case representation for inert systems

The CBR approach aims at reusing experiences in improving the settings for the inert system. A case records the experience in corrective actions to keep a target value within assigned limits around a setpoint value, for instance the room temperature within a corridor of 19.5 to 20.5 degree celsius.The case $Case = (P, S, A)$ comprises a problem description P, a solution description S and a quality assessment A of the proposed solution.

- P – The problem description records the state of the system and its environment, including the recent settings.
- S – The solution description addresses a revision of the settings by corrective actions.
- A – The quality assessment contains the results of a post-mortem analysis of the suggested solution.

Time series for *setpoint values*, *measured process values*, and *disturbance values* are recorded for the problem description P. Setpoint values describe the desired state of the inert system. Measured process values are the actual values that might deviate from the setpoint values. Disturbance values are values that have an impact on the measured process variables in addition to the corrective actions. From the point of view of the controller system, they "disturb" the control processes. From the point of view of the users, they are key determinants on the inert system. The values for the time series are recorded at equidistant time points $t_{-m}, ..., t_{-2}, t_{-1}, t_0$ with t_0 denoting the current time point at reasoning time. The continuation of some of them might be predicted for the equidistant time points $t_1, t_2, ..., t_n$, estimated at time point t_0. An example for the latter is wheather forecasting data.

The solution S records *corrective actions* that are taken to keep the measured process values within a corridor of values around the setpoint values. For the sake of simplicity, we have chosen that the solution is a single corrective action for the setting of the inert system in the next time step initially. We assume that the time interval until the next time point is large enough to measure a first impact of the corrective action. Alternatively, the solution can be described as time series of corrective actions over a number of time steps. Even interleaved phases of retrieve and reuse are possible in principle. However, the latter would lead to concurrent processes which are difficult to handle.

Table 1 illustrates a sample case in an EMS. The distance between time points is one hour. Disturbance values are the solar radiation in minutes per hour (Sun) and the temperature outside the building ($\mathrm{T}_{outside}$). The measured process values are the measured room temperature (T_{inside}). The setpoint values are the desired room temperature (T_{target}). The corrective action is the energy (E_{in}) supplied or dissipated via the EMS during the next time step t_1. The case has been recorded at time point t_0. The disturbance values have been measured until t_0. The values for t_1 until t_n are forecast data.

The assessment A is taken when the time frame is over. When time point t_n has passed, the updated values for $t_1, ..., t_n$ are used to assess the case. The predicted disturbance values have been replaced by the measured disturbance values. The time series for the measured process values and the corrective actions have been continued. The assessment considers the deviation between setpoint and measured process values by an error function e as well as the corrective actions by an energy consumption function u for $t_1, ..., t_n$. It is computed by the assessment function f for a case c by a weighted sum as follows:

$$f(c) = \sum_{i=1}^{n} w_1 \cdot e(i) + w_2 \cdot u(i)$$

The error function e measures the deviation of the actual room temperature T_{inside} from the setpoint value T_{target}:

$$e(i) = |T_{inside}(i) - T_{target}(i)|$$

The energy consumption function u measures the heating or cooling energy of the corrective action. Since the production of cooling energy consumes nearly twice the energy of heating [7] , we multiply cooling energy with the factor 2:

$$u(i) = \begin{cases} E_{in}(i), & E_{in}(i) \geq 0 \\ E_{in}(i) \cdot 2, & E_{in}(i) < 0 \end{cases}$$

The weights w_1 and w_2 specify the balance between reaching the target temperature and saving energy.

4 Case retrieval

A time event triggers a reasoning cycle starting with the retrieve phase. We have chosen hourly time events. The query describes the current situation of the EMS,

Time point	**Time stamp** [$yyyymmddhh$]	**Sun** [$\frac{min}{h}$]	$\mathbf{T}_{outside}$ [$°C$]	$\mathbf{T}_{inside}$ [$°C$]	$\mathbf{T}_{target}$ [$°C$]	$\mathbf{E}_{in}$ [Wh]
t_{-m-2}	1981083016	60	20.3	19.94	20	-3500
t_{-m-1}	1981083017	42	20.1	20.06	20	-750
t_{-m}	1981083018	0	18.8	21.43	20	0
...	1981083019	0	16.9	21.12	20	0
	1981083020	0	15.4	20.63	20	0
	1981083021	0	13.8	20.02	20	1000
	1981083022	0	13.2	20.07	20	1000
	1981083023	0	11.1	20.04	20	1250
	1981083100	0	10.4	19.96	20	1500
	1981083101	0	10.1	20.00	20	1500
	1981083102	0	9.1	20.00	20	1750
...	1981083103	0	8.0	20.08	20	1750
t_{-2}	1981083104	0	8.1	20.02	20	1750
t_{-1}	1981083105	0	8.3	19.98	20	0
t_0	1981083106	48	9.1	18.62	20	0
t_1	1981083107	60	11.5	–	20	**?**
...	1981083108	60	14.0	–	20	–
	1981083109	60	16.0	–	20	–
	1981083110	60	17.8	–	20	–
	1981083111	60	19.1	–	20	–
	1981083112	54	20.0	–	20	–
	1981083113	42	20.3	–	20	–
	1981083114	60	20.9	–	20	–
	1981083115	42	21.4	–	20	–
	1981083116	0	21.2	–	20	–
...	1981083117	0	21.0	–	20	–
t_n	1981083118	0	19.7	–	20	–
t_{n+1}	1981083119	0	18.2	–	20	–
t_{n+2}	1981083120	0	17.1	–	20	–

Table 1. The problem description of a sample case.

including the setpoint values, measured process values, and disturbance values. The case depicted in Table 1 can serve as a sample query. The retrieval uses a composite similarity measure for a query and a case that aggregates the local similarity measures by a function F:

$$
\begin{aligned}
sim = F(&sim_{Time\_stamp}, \\
&+ sim_{Sun}, \\
&+ sim_{T\_outside}, \\
&+ sim_{T\_inside}, \\
&+ sim_{T\_target})
\end{aligned}
$$

F is a weighted sum. sim_{Time_stamp} considers the annual date $date$ and the time of day $hour$ when the query and the case were recorded each:

$$
sim_{Time\_stamp}(query, case) = \frac{1}{1 + |date_{query}(t_0) - date_{case}(t_0)| \cdot |hour_{query}(t_0) - hour_{case}(t_0)|}
$$

We assume the values of the time series Sun, $T_{outside}$, T_{inside}, and T_{target} as vectors. The local similarity measures for the time series are computed by means of the City Block Metric [1]. The size of the vectors Sun, $T_{outside}$, and T_{target} is $m + n + 1$ and $m + 1$ for T_{inside}, since T_{inside} data only exists for the past.

As a starting point, we have chosen straight forward similarity measures. We will investigate further, more sophisticated similarity measures for time series, such as dynamic time warping [10], as a part of our future work.

5 Case reuse

The solution of the best matching case is reused for the current situation. The solution describes the corrective action for the settings of the system by the amount of energy $E_{in\_case}(t_1)$ to be infused into the building next. However, E_{in_case} has to be adapted to the recent situation. The impact of an energy infusion depends not only on the bare amount of energy supplied or distracted but also on the current room temperature and on the heat capacity of the building. The latter can be specified by a constant $c_{building}$.

The amount of energy is adapted as follows:

$$
e_{ad} = E_{in\_case}(t_1) + c_{building} \cdot (T_{inside\_case}(t_0) - T_{inside\_query}(t_0))
$$

The difference between the room temperature of the reference case and the current temperature in the building is multiplied with $c_{building}$. In case the current room temperature is lower, more energy is required for heating than in the case (or the energy that is required for cooling can be reduced, i.e. the negative value of E_{in_case} increases). In case the current room temperature is higher than in the case, E_{in_case} decreases analogously. The adaptation could lead to amounts of energy that are not available for heating and cooling in our building. Thus, we introduce the limits E_{min} for cooling and E_{max} for heating. The final amount of energy to be infused is determined by the following clipping function:

$$E_{ad}(t_1) = \begin{cases} E_{min}, \ e_{ad} < E_{min} \\ e_{ad}, \quad E_{min} < e_{ad} < E_{max} \\ E_{max}, \ else \end{cases}$$

In future work, the approach might be extended to reuse a sequence of corrective actions $E_{in\_case}(t_1)...E_{in\_case}(t_k)$.

6 Experimental setup

We have implemented the case-based approach for inert systems and conducted an experimental evaluation with an EMS scenario. The results of the CBR system have been compared to a traditional PID controller with respect to energy consumption and comfort for the occupants.

Ideally, the experiments would be executed in a real building measuring the energy consumption by sensors at the valves and measuring the comfort by acquiring feedback from the real occupants. Since these resources are difficult to obtain, the experiments have been conducted in silico. They involve an energy simulation of a building to approximate the impact of both, the energy infusion by the system as a corrective action and the two metereological parameters sun duration and outside temperature as disturbance variables. A seeding case base has been constructed from real weather data for the time period from 1981 to 2014. The experiments on the behavior of the CBR system were then simulated with the weather data for the year 2015.

In our example we use a grid of one hour for all of our calculations. On the one hand this reflects the inertness of a building. On the other hand this decision is taken to limit the computational complexity.

The energy model of the building assumes a single cubic room with an edge length of 10m that has one side with glass windows. The relative position of sun is not taken into account. Basically, the temperature of the air in a building depends on the energy flow into and out of the building through walls and windows. The loss or gain of the energy ΔE through walls and windows can be calculated by using the thermal transmittance (also known as U-value or k-value):

$$\Delta E = U \cdot \Delta T \cdot A \cdot \Delta t$$

U characterizes the isolation value of the wall or the window. ΔT is the temperature difference between $T_{outside}$ and T_{inside}. A is the surface that divides inside and outside and Δt is the time span.

The dynamic simulation is done by the iteration:

$$E_{t+1} = E_t + \sum_{i=1}^{n} \Delta E_n$$

where ΔE_n represents the different sources and drains of energy. So far we use ΔE_1 for the energy flow through walls, ΔE_2 for the energy flow through windows, ΔE_3 for the energy of the sun through the windows and ΔE_4 for the energy of the heating and cooling system.

As the base for the dynamic simulation of the building we use historical weather data of Frankfurt a.M./Germany ranging back to 1981 in an one hour resolution containing air temperature, air speed and direction, humidity and minutes of sun per hour. It is important to note that the initial case base is built for that climate and applies only to regions within the same climate classification.

In a post-mortem analysis, we created a case for each hour of the historical weather data as t_0. The setpoint values are fixed to 20.0°C. We calculated optimal energy values to be infused into the building for four hours, i.e. $t_n = t_4$ regarding the development of the weather and the (simulated) state of the building. In a brute-force approach, we explored the full solution space with energy amounts between $E_{min} = -4kWh$ and $E_{max} = 4kWh$ in a grid of granularity $g = 0.1kWh$. The assessment function f (compare Section 3) serves as fitness function to optimize the E_{in} values. For our example we weighted the energy consumption function with zero ($w_2 = 0$) to force the system to generate seeding cases that keep the given temperature as good as possible. The complexity of generating an optimum initial case follows $\mathcal{O}((\frac{E_{max}-E_{min}}{g})^n)$.

Instead of using the simple approach for creating a seeding case base as described above, a wide variety of modifications for a real building is possible and desirable. Changing demands on the target temperature regarding the comfort of the inhabitants of the building are one example. Another opportunity for an extension is to use sliding frames of acceptable min/max temperatures to preserve a maximum of energy.

We designed two variants of an experiment to evaluate our system. The first experiment tests the ability of our system to compete with a common PID controller if it has access to future weather data. The CBR system uses the future weather data to find the best matching case. The similarity of two cases is calculated with a time span of ±12 hours where we assume the existence of a high-quality weather prediction for the next 12 hours. The second experiment explores the behavior of our system if no future weather data is available. It acts on the same seeding case base as the first experiment, but the best matching case must be retrieved without knowledge of the future weather development. The similarity of two cases is based on the data for the previous 24 hours.

We compared our results against a hand optimized PID controller. To keep the computational effort low, this PID controller works as all of our calculations in the one hour grid that we use for the CBR system. On the one hand this decision is debatable, since a real PID controller works in a grid of seconds or minutes and can thus adapt much faster to changes of the impact factors. On the other hand, the delay between demand and delivery of hot and cold water (or air) in a real building lies between 15 minutes an 30 minutes. Arguably, our one hour computation grid is not as precise as a real system but it acts similarly.

The PID controller uses the difference between inside and outside temperature as the base for its calculations.

7 Experimental results

For both experiments, we measured the comfort and the entire energy consumption for the year 2015, comparing the PID controller and CBR system. The comfort is measured by the Root Mean Square Error (RMSE) of the deviations from the desired room temperature of 20°C within the simulated building. Second, we measured the entire energy consumption for the year 2015 at large.

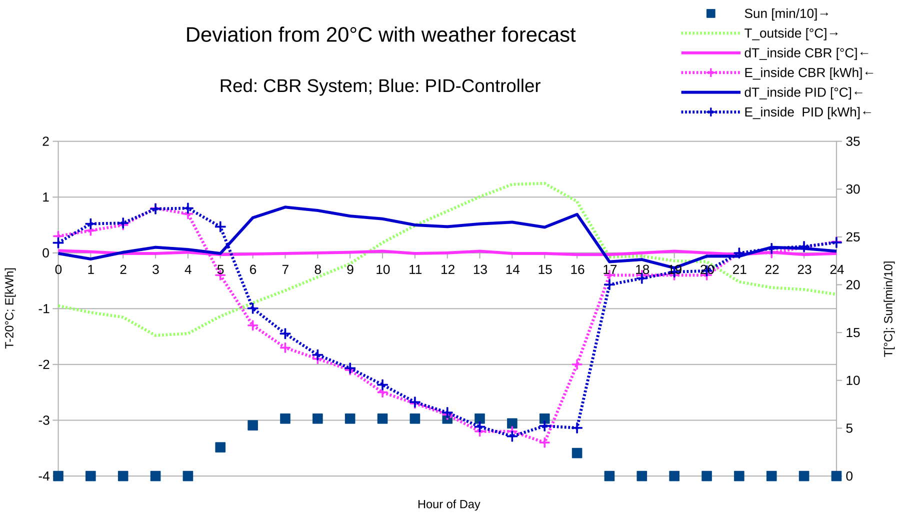

Fig. 1. Comparison of CBR (bright red) and PID (dark blue) with the objective to maintain exactly 20°C. Future weather data is available.

For our first experiment, that considers future weather data, an interesting example is the situation for the 12th of June 2015 as depicted in Figure 1. The outside temperature increases until the early afternoon where the sun seems to be hidden by clouds. During the sunny period, the PID controller results in a room temperature that is slightly higher than 20.5°C which is the upper boundary of the comfort corridor. The CBR system maintains the 20°C nearly perfectly. The RMSE of the deviation for the whole year 2015 of the temperature inside the building is 0.035°K for our CBR system and 0.32°K for the PID controller. A more palpable metric is the added up deviation of the temperature: $\sum_{2015010100}^{2015123123} |20.0°C - T|$ which amounts to 178°K for the CBR system and 1595°K for the PID controller for the entire year.

The energy infusion values depicted in the sample in Figure 1 seem very similar for both, the CBR system and the PID controller. The cumulative annual values confirm this observation. The CBR system used 13.4MWh of energy for heating and cooling of the building. The PID used 13.7MWh of energy. This surprising coincidence can be explained simply by the fact that the PID controller infuses too much energy (overprovisioning) about as frequently as too little energy (underprovisioning).

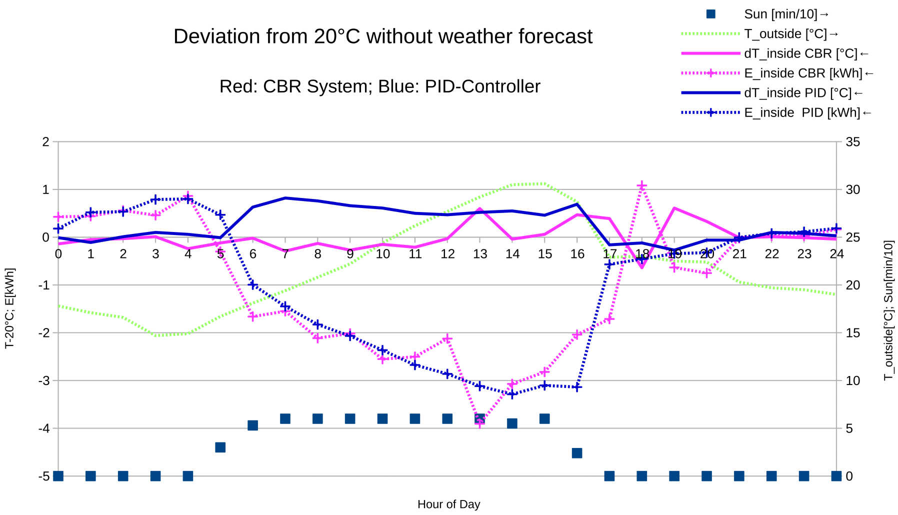

Fig. 2. Comparison of CBR (bright red) and PID (dark blue) with the objective to maintain exactly 20°C. Future weather data is not available.

The second experiment without considering future weather data leads as expected to slightly worse (but still reasonable) results for the CBR system. Obviously, the values for the PID controller remain the same as in the first experiment. The results for the same sample day as above is depicted in Figure 2. It can be seen that both curves for room temperature and energy infusion by the CBR system are less smooth than in the first experiment. The annual RMSE for the CBR system amounts to 0.26°K. The added up deviation for the cbr system is 1505°K. This value is still slightly better than the 1595°K for the PID controller for the entire year. The CBR system used 13.8MWh of energy for heating and cooling of the building. This is slightly worse than the 13.7MWh of the PID.

8 Acknowledgements

The authors would like to thank Prof. Alfred Karbach and his team at the THM University of Applied Sciences for providing valuable insights to energy systems and many fruitful discussions on user comfort. Further, we thank Dipl.-Ing. Beate Massa whose expert advice on facility management issues we appreciate very much.

9 Discussion and conclusion

We have introduced a novel CBR approach for the experience-based control of inert systems and demonstrated the feasibility of the approach in the field of energy management.

We have presented a case representation with time series of impact factors and state descriptors, including setpoint values, measured process values, disturbance values and corrective actions. A straightforward similarity measure has been specified. An adaptation rule considering physical properties such as the heat capacity of a building has been proposed. Our experimental results for a simulated building under real weather conditions provide a proof of concept for using CBR for building EMSs. The experimental results are quite promising in comparison to a traditional PID controller. The first experiment compares the CBR system with a common PID controller if CBR has access to future weather data. The second experiment compares the systems if there is no such access. The first experiment has clearly shown that CBR outperforms the traditional PID in terms of both, energy consumption and comfort. The second experiment has shown that CBR is competitive to traditional PID by comparable values for energy consumption and by better comfort values.

In contrast to PID, CBR provides a wide range of opportunities for further improvements. In addition to weather forecasts, further aspects of forward planning might be considered, such as expected occupancy, or renewable energy availability. This would allow us to extend the quality function to cost aspects and, hopefully, to save both, energy and money.

The next step of our future work will be to create an experimental setup in-vitro. Thus, we are planning to confirm the simulation results from the in-silico experiments by measured values to gain further experiences with the system, for instance on the optimal length of time intervals.

Further intriguing research questions are whether the cases can be transferred to other buildings, or whether the approach can be tranferred to other application scenarios for inert systems, such as in medicine. We believe that CBR is capable of providing significant benefits for the control of inert systems, especially in reducing modelling efforts and energy consumption.

References

1. Bergmann, R.: Experience management: Foundations, development methodology, and Internet-based applications. Springer Verlag (2002)

2. Gundersen, O.E.: An Analysis of Long Term Dependence and Case-Based Reasoning. In: Workshop Proceedings of ICCBR 2014. pp. 129 – 138 (2014)
3. Gundersen, O.E., Bach, K.: RATIC 2016: Workshop on Reasoning about Time in CBR. In: Workshop Proceedings of ICCBR 2016 (2016)
4. Gundersen, O.E., Montani, S.: RATIC 2014: Workshop on Reasoning about Time in CBR. In: Workshop Proceedings of ICCBR 2014, pp. 125 – 164 (2014)
5. Leake, D., Wilson, D.: Taking Time to Make Time: Types of Temporality and Directions for Temporal Context in CBR. In: Workshop Proceedings of ICCBR 2014. pp. 149 – 153 (2014)
6. Ma, Y.: Model Predictive Control for Energy Efficient Buildings. Ph.D. thesis, UC Berkeley (2012)
7. Massa, B.: Personal interview on energy cost calculations in facility management, 2017-04-25 at Frankfurt a.M., Germany (2017)
8. Nise, N.S.: Control Systems Engineering, 6th Edition. Wiley (2011)
9. Platon, R., Martel, J., Zoghlami, K.: CBR Model for Predicting a Building's Electricity Use: On-Line Implementation in the Absence of Historical Data. In: Case-Based Reasoning Research and Development - 23rd International Conference, ICCBR 2015, Frankfurt am Main, Germany, September 28-30, 2015, Proceedings. pp. 306–319. LNCS 9343, Springer (2015)
10. Richter, M.M., Weber, R.: Case-Based Reasoning: A Textbook. Springer (2013)
11. Shabani, A., Paul, A., Platon, R., Hüllermeyer, E.: Predicting the Electricity Consumption of Buildings: An Improved CBR Approach. In: Case-Based Reasoning Research and Development - 24th International Conference, ICCBR 2016, Atlanta, GA, USA, October 31 - November 2, 2016, Proceedings. pp. 356–369. LNCS 9969, Springer (2016)