2013 European Control Conference (ECC)
July 17-19, 2013, Zürich, Switzerland.

Learning/Repetitive Control for Building Systems with Nearly Periodic Disturbances

Kasper Vinther, Vikas Chandan, and Andrew G. Alleyne

Abstract— **In this paper, learning/repetitive control is proposed for improvement of existing feedback control loops for temperature regulation in buildings. A single zone office building is used as an example, with real weather data for Phoenix Arizona and realistic occupancy load schedules. Simulations have shown a decrease in the average set point tracking error of more than 50%, even without additional energy consumption. This can be achieved in situations where the load disturbances have enough repeatability and a repeatable-to-nonrepeatable ratio can be computed to determine if learning should be used and at which frequencies. Furthermore, the increased tightness in reference tracking could be used to lower energy consumption by moving the reference set point closer to the boundaries of the allowable temperature range.**

I. Introduction

Temperature control in buildings is important for various reasons. It could be used to improve occupancy comfort in commercial and residential spaces, to maintain good quality of products in cold storage rooms or for safety reasons in pharmaceutical processing facilities. The temperature is often controlled with a heating ventilation and air-conditioning (HVAC) system using either on/off type hysteresis control or variations of PI/PID control due to their simple implementations [1], [2]. However, they seldom perform well at all possible operating conditions because the control is typically tuned only for a selected nominal situation; also the control might not be tuned properly in the first place [2], [3], [4].

Model based control can improve the performance over on-off or PID control, but obtaining a sufficiently accurate model of the building and the HVAC system is a time consuming and costly process and the uncertainty of these models can be quite severe due to several reasons. Firstly, certain parameters are expected to be time varying, such as the thermal capacities, and therefore are hard to obtain. Secondly, there will be uncertainty in the estimation of load disturbances, which are dictated by several factors such as weather conditions, occupancy, appliances, lighting, etc. Although weather can be included in the control [5], the exact contribution of factors such as occupants, appliances, and thermal infiltration is difficult to predict accurately.

This paper investigates the use of learning type control such as iterative learning control (ILC) and repetitive control (RC) for improvement of existing feedback controllers in buildings such as the ones mentioned above. A survey on ILC is provided in [6] and RC is covered in references such as [7], [8], which also show the similarity of these methods. The basic concept of these types of learning control methodologies is that feedforward signals are generated and updated based on previous errors and they do not require the costly process of identifying system models. However, they do need repetitiveness in the tracking error and disturbances.

Buildings experience repetitiveness in load disturbances on a daily and a seasonal basis due to the weather. A good prediction of the weather tomorrow is that it will most likely be close to what it was today. Also, the occupancy and appliance loads will often repeat themselves on a daily basis governed by actions such as people going to work, having lunch, supermarkets opening and closing, etc. The assumptions about the repetitiveness of various disturbances will not always hold true, but the idea in learning is that the control adapts to changes based on past experience. This means that we only need nearly periodic disturbances, in the sense that they are allowed to change, but on average they will show periodicity. The authors in [9] have provided a way to calculate the Repeatable-to-Nonrepeatable Ratio (RNR), which can be used to determine if the repeatable part of the tracking error or disturbance is larger than the nonrepeatable part for each frequency. A ratio above zero means that there is potential for applying learning type control.

A single zone office building is provided as an example case study to show the possible tracking performance improvement with learning/repetitive control. A MATLAB Simulink model of this building is derived using lumped parameter modeling with resistive-capacitive networks and quasi periodic disturbances such as real weather data and occupancy load schedules are used. This model allows us to simulate an entire year or even longer and makes it possible to compare different control strategies, which would not be possible in a real building. The proposed steps to improve tracking performance in an already existing stable feedback system are to: (i) Collect error data for an appropriate time window. (ii) Calculate RNR. (iii) Use RC or ILC to design a learning filter that corresponds to the RNR findings and thus improve tracking performance at appropriate frequencies.

This paper first introduces a simple modeling framework for buildings and provides a description of the model of a single zone office building. A temperature controller is then designed in Section III, which provides a comparative case for analysis. Then, the learning/repetitive control design is described in Section IV, followed by results in Section V, from simulations on the single zone building spanning an entire year. Finally, conclusions are drawn in Section VI.

K. Vinther kv@es.aau.dk, Section of Automation and Control, Department of Electronic Systems, Aalborg University, 9220, Denmark.

V. Chandan vchanda2@illinois.edu and A. G. Alleyne alleyne@illinois.edu, Department of Mechanical Science and Engineering, University of Illinois, Urbana, IL - 61801, USA.

II. Modeling

A. Lumped Parameter Modeling of Buildings Using Resistive-Capacitive Networks

A commonly used method of modeling the thermal dynamics of buildings is using lumped resistive-capacitive networks [10], [11], [12]. In this paper, each room (zone) is considered as a single capacitance and each wall is represented with 3 resistors and 2 capacitors also referred to as the 3R2C modeling approach [11]. Fig. 1 illustrates the 3R2C network, where $R_{wo,i}$ is the thermal resistance between the outer wall surface of wall i and ambient, $R_{wi,i}$ is the resistance between the inner wall surface and ambient, and $R_{ww,i}$ is the resistance between the inner and outer wall surfaces. The states in the model are the temperatures of the

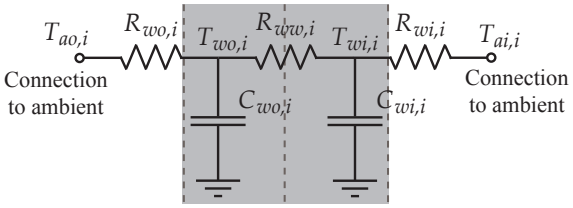

Fig. 1. 3R2C network model of a wall.

inner wall surfaces $T_{wo,i}$, the outer wall surfaces $T_{wi,i}$, and the room/zone temperatures. The ambient temperatures $T_{ao,i}$ and $T_{ai,i}$ could be zone temperature, outside air temperature, or ground temperature, where outside air (T_{air}) and ground (T_{gnd}) temperatures are considered inputs to the system.

The zones have controllable inputs, which are the heat transfer rates $\dot{\mathbf{Q}}_{hvac}$ supplied by the HVAC system (positive when heating and negative when cooling). A discrete state space model representation of this system is shown in (1).

$$\begin{bmatrix}\mathbf{T}_w\\ \mathbf{T}_z\end{bmatrix}(k+1)=\begin{bmatrix}\mathbf{A}_{ww} & \mathbf{A}_{wz}\\ \mathbf{A}_{zw} & \mathbf{A}_{zz}\end{bmatrix}\begin{bmatrix}\mathbf{T}_w\\ \mathbf{T}_z\end{bmatrix}(k)+\begin{bmatrix}\mathbf{0}\\ \mathbf{B}_z\end{bmatrix}\dot{\mathbf{Q}}_{hvac}(k)+\begin{bmatrix}\mathbf{B}_{air} & \mathbf{B}_{gnd} & \mathbf{B}_{dw} & \mathbf{0}\\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{B}_{dz}\end{bmatrix}\begin{bmatrix}T_{air}\\ T_{gnd}\\ \mathbf{d}_w\\ \mathbf{d}_z\end{bmatrix}(k) \quad (1)$$

The vector $\mathbf{T}_w$ contains all the wall surface temperatures (both inside and outside) and the vector $\mathbf{T}_z$ contains all the zone temperatures. T_{air}, T_{gnd}, $\mathbf{d}_w$, and $\mathbf{d}_z$ are considered as disturbances to the system, where $\mathbf{d}_w$ is a vector of lumped Long Wave Radiation (LWR) and Short Wave Radiation (SWR) heat transfers affecting each wall and $\mathbf{d}_z$ is a vector of the thermal load applied directly on each zone (e.g. occupants, appliances, and lighting). These are further described in Subsection II-B.

B. Single Zone Office Model

In this paper, we use a single zone office building with four walls, a ceiling, and a floor as an example, which is assumed to be located in Phoenix, Arizona. Using a standard medium office building construction template provided in the OpenStudio tool [13] developed by the National Renewable Energy Laboratory (NREL) an EnergyPlus [14] model of this building was created, which uses a weather file for Phoenix, Arizona. A Google SketchUp illustration of the building is shown in Fig. 2.

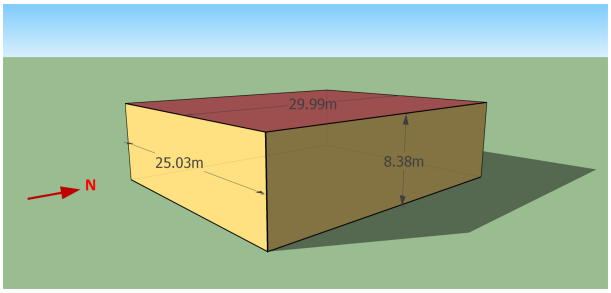

Fig. 2. Google SketchUp illustration of single zone EnergyPlus model created in OpenStudio.

The 3R2C framework described in Subsection II-A was used to construct a simplified resistive-capacitive network model in MATLAB. This model is based on the EnergyPlus model, where the resistances corresponding to the heat transfer between the zone and wall surfaces, and the zone capacitance were obtained using a least squares based system identification procedure applied on data generated using EnergyPlus. The capacitances for the inner and outer wall surfaces, and the resistances between them were computed based on physical properties of the wall construction layers. The resistances corresponding to thermal interactions of the wall surfaces with the ground and ambient air were obtained by averaging the corresponding heat transfer coefficient data obtained from a yearlong EnergyPlus simulation. The floor is modeled with only one state $T_{wi,6}$ as the ground temperature under the floor ($T_{wo,6}$) is considered constant.

The thermal load disturbances acting on the eleven wall surfaces are collected in a single vector $\mathbf{d}_w(t) = [d_{wi,1}(t), d_{wo,1}(t), ..., d_{wi,5}(t), d_{wo,5}(t), d_{wi,6}(t)]^T$ consisting of the disturbance on the inner wall surfaces ($i \in 1, 2, ..., 6$)

$$d_{wi,i}(t) = \underbrace{\sigma A_i \sum_{j=1}^{6} \mathbf{F}_{i,j}\left(T^4_{wi,j}(t) - T^4_{wi,i}(t)\right)}_{LWR\ from\ wall\ surfaces}, \quad (2)$$

and the outer wall surfaces ($i \in 1, 2, ..., 5$)

$$\begin{aligned} d_{wo,i}(t) = {} & \underbrace{\epsilon_i \sigma A_i F_{gnd,i}\left(T^4_{gnd}(t) - T^4_{wo,i}(t)\right)}_{LWR\ from\ ground} \\ & + \underbrace{\epsilon_i \sigma A_i F_{sky,i}\left(T^4_{sky}(t) - T^4_{wo,i}(t)\right)}_{LWR\ from\ sky} \\ & + \underbrace{\epsilon_i \sigma A_i F_{air,i}\left(T^4_{air}(t) - T^4_{wo,i}(t)\right)}_{LWR\ from\ air} + \underbrace{\alpha_i A_i q_{SWR,i}(t)}_{SWR}, \end{aligned} \quad (3)$$

where i is the wall number, σ is the Stefan-Boltzmann constant, A_i is the surface area, $\mathbf{F} \in \mathbb{R}^{6x6}$ is a matrix of Script-F factors [15], $F_{gnd,i}$, $F_{sky,i}$, and $F_{air,i}$ are view factors for the outer surface of wall i, T_{sky} is the sky

temperature, ϵ_i is the thermal absorbtance of wall i, α_i is the solar absorbtance of wall i, and q_{SWR} is the incident solar radiation per unit area on wall i. These values, together with T_{air} and T_{gnd}, can be obtained directly from EnergyPlus for an entire year and (2) and (3) emulate how EnergyPlus calculates the disturbances (short and long wave radiation).

The zone load d_z is pseudo randomly generated for each day for an entire year. Each day we assume that five Gaussian probability density functions (PDF) govern how many people arrive for work, how many hours each person works, when each person work (midpoint of working hours), when each person has lunch (midpoint), and how long that person's lunch is. The PDF parameters are mean $\bar{x} = \{45$ people, 9 hours, 13:00 hours, 12:30 hours, 0.75 hours$\}$ and standard deviation $\sigma_{std} = \{3$ people, 1 hour, 0.5 hours, 0.25 hours, 0.1 hours$\}$. Each occupant in the office corresponds to 0.6 kW (0.1 kW body heat, 0.4 kW appliances, 0.1 kW lighting) and 0.3 kW when they are at lunch (0.2 kW appliances, 0.1 kW lighting). For simplicity there is no distinction of weekends, however, one could just turn the learning off during weekends or even have different learning controllers for weekdays and weekends.

In this case study, the controllable heat transfer rate to the zone $\dot{Q}_{hvac}$ is assumed to be bounded between -35 and 17.5 kW. This will result in controller saturation during some hot summer days and thus requirement for integrator anti-windup. Furthermore, the zone temperature reference $T_{z,r}$ is set to 22.5°C. A reference that varies during the day could also be used, as long as it repeats itself daily.

III. Office Building Temperature Control

The air temperature in office buildings can be controlled in different ways, e.g. with simple on/off hysteresis based control, PI/PID control, or more advanced model based control. In most cases one of the first two methods is used [1], [2], [3], [4], primarily due to their simple implementations and relatively simple tuning procedures. Fig. 3 shows the closed loop feedback system, where the zone air temperature T_z tracks a reference temperature $T_{z,r}$. The system is denoted

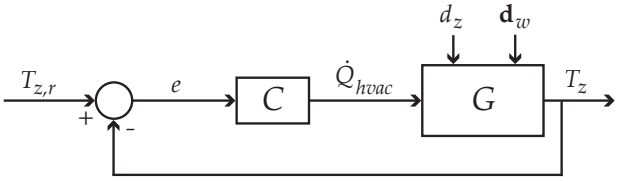

Fig. 3. Closed loop feedback control structure for temperature control.

by G and has the heat transfer rate provided by an HVAC system, $\dot{Q}_{hvac}$, as input. The disturbances to the system are the thermal loads on the walls $\mathbf{d}_w$ and the thermal load on the zone air d_z. The system model and disturbances in the single zone building case study are defined in Subsection II-B and are assumed to be unknown to the feedback controller.

Sensing of the zone air temperature with a digital sensor (thermostat) and its communication with the controller C is performed in discrete time. A realistic sample time for buildings is somewhere between 1 to 15 minutes depending on the size of the zones among other factors. Here a 5 minute sample time t_s is used. This sampling rate is faster than that of the underlying system dynamics and slow enough to allow the HVAC system to meet the desired $\dot{Q}_{hvac}$ before the demand is changed again. For comparison, the open loop zone temperature time constant is about 27 minutes and the wall surface temperature time constants are much larger.

PI control is used in the feedback loop in this case study, since it performs better than just having hysteresis based control and because it is commonly used as mentioned earlier. There exist various different PI tuning methods ranging from step response analysis, relay feedback, robust synthesis, etc. However, robust control methods require a nominal model of the system and uncertainty models, which can be difficult to obtain as mentioned earlier, and are usually not available to the control engineer. The step response method is therefore proposed, which gives different results depending on the particular load on the system. Critical regions of operation correspond to high and low load situations, which should be considered in order to design a robust PI controller. Fig. 4 shows closed loop step responses using three different discrete-time test controllers C_1, C_2, and C_3 (4) for a high load situation (Fig. 4(a) and 4(b)) and a low load situation (Fig. 4(c) and 4(d)). The high load situation is generated

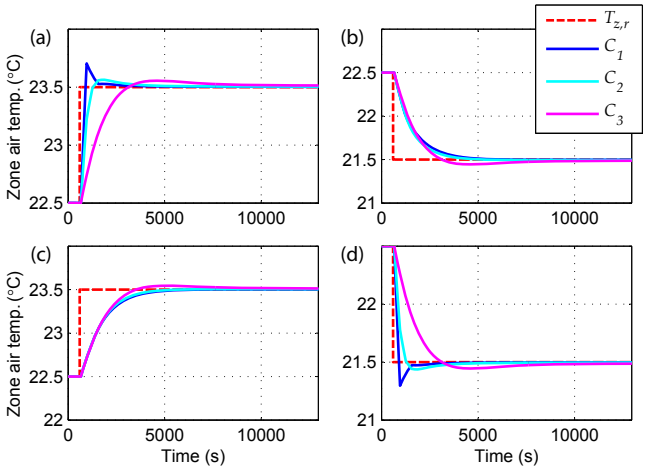

Fig. 4. Closed loop step response with different controllers C_n during high load (a,b) and low load (c,d).

based on the air and sky temperature at 1 pm July 1 with zone load $d_z = 25$ kW and the low load situation is generated based on conditions at midnight on January 1 and with zone load $d_z = 0$ kW.

$$C_n = K_{p,n} + K_{i,n} t_s \frac{z}{z-1}, \; n \in \{1,2,3\}, \qquad (4)$$

$$\begin{aligned} \text{where } & K_{p,1} = 25, K_{i,1} = 0.025, \\ & K_{p,2} = 15, K_{i,2} = 0.015, \\ & K_{p,3} = 5, K_{i,3} = 0.005. \end{aligned}$$

The step up in reference at high load and the step down at low load determines how fast a controller we can design. In

these situations, the controller C_1 is too aggressive, giving a large overshoot in just one time step. Controller C_3 on the other hand is too conservative, whereas C_2 provides a good balance with a maximum overshoot of 6% and a minimum rise time of approximately one time step (300 seconds). This corresponds to Fig. 4(a) and 4(d). The slowest responses are observed in Fig. 4(b) and 4(c), where the rise time is between 900-1200 seconds. This confirms that a fixed gain PI control is not ideal under all operating conditions and therefore motivates the development of a learning/repetitive type control for improvement of tracking performance.

IV. Learning/Repetitive Control for Improved Reference Tracking

Fig. 5 illustrates the proposed addition of RC to the feedback control loop shown in Fig. 3. The RC block

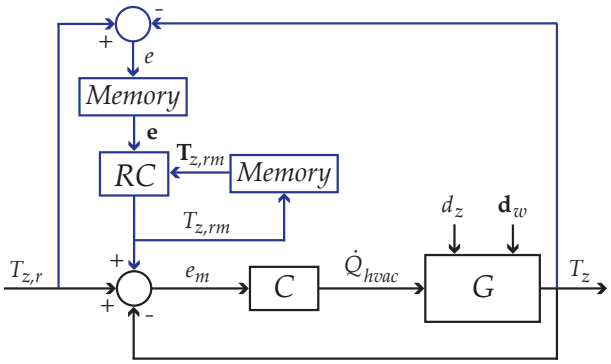

Fig. 5. Existing feedback loop (black) with additional RC (blue).

modifies the original temperature reference $T_{z,r}$ based on the tracking error one trial back in time e (one trial is a repetitive cycle, e.g. 24 hours), which will lower e as the number of trials goes to infinity, if the load pattern is repetitive enough.

Collecting past errors e and reference modifiers $T_{z,rm}$ in vectors, $\mathbf{T}_{z,rm} = [T_{z,rm}(k-2N+1), ..., T_{z,rm}(k-2), T_{z,rm}(k-1)]$ and $\mathbf{e} = [e(k-2N+2), ..., e(k-1), e(k)]$, where k is the current time step, and using buffers (memory blocks in Fig. 5) of size $2N-1$, where N is the number of discrete samples in one repetitive trial, we can define a repetitive controller as

$$\hat{\mathbf{T}}_{z,rm} = Q(\mathbf{T}_{z,rm} + L\mathbf{e}). \tag{5}$$

The vector $\hat{\mathbf{T}}_{z,rm} = [\hat{T}_{z,rm}(1), \hat{T}_{z,rm}(2), ..., \hat{T}_{z,rm}(2N-1)]$ is the output of the learning algorithm and it contains updated reference modifiers. However, we are only interested in $\hat{T}_{z,rm}(N+1)$ since this is the updated version of the reference modifier during the last trial (N samples back in time) to be applied in the next time step $k+1$. Therefore, the reference modifier for the next time step is

$$T_{z,rm}(k+1) = \hat{T}_{z,rm}(N+1). \tag{6}$$

The calculations are repeated at each time step and L and Q can be causal/non-causal filters with low-pass, band-pass or high-pass characteristics depending on the frequencies of the repetitive part of the error. Choosing L and Q as filters is also the reason why vectors are used in (5).

The proposed RC is similar to a serial structure ILC implementation, with the only difference being that the reference modifier is calculated in real time at each time step k, whereas in ILC, all reference modifiers are calculated after a trial has ended for the entire subsequent trial. The benefit of RC is that it does not require the same initial condition each trial. For more detail on discrete RC see [7], [16], [17].

In order to determine the cutoff frequencies for the learning filter L we propose to calculate the Repeatable-to-Nonrepeatable Ratio (RNR) for all frequencies. This method is shown in [9] and calculated as

$$RNR = 20log\left(\frac{|FFT[\bar{\mathbf{e}}]|^2}{\frac{1}{N_t}\sum_{j=1}^{N_t}|FFT[\bar{\mathbf{e}}-\mathbf{e}_j]|^2}\right), \tag{7}$$

where N_t is the number of trials used in the analysis, FFT is the Fast Fourier Transform, and the repeatable error is

$$\bar{\mathbf{e}} = \frac{1}{N_t}\sum_{j=1}^{N_t}\mathbf{e}_j, \tag{8}$$

where $\mathbf{e}_j$ is a vector of all the errors in trial j. Equation (7) calculates the power of the repeatable signal versus the power of the nonrepeatable signal and converts it to dB. If this number is larger than 0 dB for a particular frequency it means that there is more repetitiveness in the error than nonrepetitiveness. The learning filter L should then be applied only on these frequencies, which can be achieved by picking appropriate cutoff frequencies for a band-pass filter or an upper cutoff frequency for a low pass filter. The filter Q can then e.g. have cutoff frequency above L in the low-pass case and is there to ensure that the condition for stability is met. Stability can be checked using (9) (derivation in [7]), which must be satisfied for all frequencies.

$$\|Q(1-zLP)\|_\infty < 1 \tag{9}$$

The transfer function P is the complementary sensitivity transfer function from reference to output defined as

$$P = \frac{GC}{1+GC}. \tag{10}$$

V. Case Study - One Year Simulation

The previously described single zone office building is used here as a case study, to show the reference tracking performance improvement of RC when applied to a system with nearly periodic load disturbance. The repetitive trial length is identified as one day starting from midnight, where the occupancy load is small.

Fig. 6(a) shows the tracking error for each day in January (blue) for the feedback control without RC and the corresponding repeatable error (red), calculated using (8), where $N_t = 30$. Furthermore, for comparison, the error for each day in an entire year is also shown (cyan). Based on the tracking error, the RNR is calculated for January and also for an entire year. However, we assume that only data from January was available as training set for the design of the

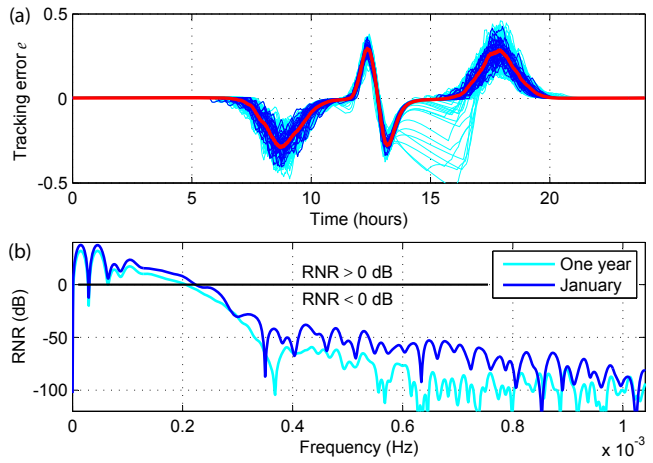

Fig. 6. The top graph (a) shows the daily tracking error e during January (blue), the repeatable error (red), and the daily error for an entire year (cyan). The bottom graph (b) shows the RNR at different frequencies based on the January errors (blue) and the entire year (cyan).

learning filter L. The uncertainty of not having each month in the year will get assigned to the nonrepeatable part of the signal. More data could be used if available.

The RNR in Fig. 6(b) indicates that L should be a low-pass filter in this particular case study and the cutoff frequency was set to 0.000223 Hz. The MATLAB command $butter$ was used to design a first order Butterworth low-pass filter and the command $filtfilt$ was then used to make a non-causal zero phase filter. The Q filter was also chosen to be a non-causal zero phase filter with a cutoff frequency of two times 0.000223 Hz. This satisfies the stability condition in (9). Additionally, the reference modifier $T_{z,rm}$ was limited between -1 and 1. This means that in worst case, if today does not look like yesterday, we will only end up following a reference of minimum 21.5 or maximum 23.5 degrees, which should still be acceptable for the occupants. Tighter constraints can be used if necessary.

After the training period, RC can be added to the system. Fig. 7 shows the infinity-norm, 2-norm, and 1-norm of the error for each day from the 1st of February to the 31st of December plotted in a histogram. The same simulation is repeated without RC for comparison. All error norms show that RC improves the reference tracking performance. This is also indicated in Fig. 8 where the 2-norm is plotted for May, June, and July. However, there are a few days where the RC performs worse (e.g. trial 152). These outliers are caused by hot summer days where the HVAC system is saturated, which is also visible in Fig. 6 with larger errors around the hot hours of the day between 1-5pm. If there is saturation one day (e.g. trial 151), the RC will learn a reference that compensates for this effect, but the next day might be colder resulting in worse tracking that following day. However, the control has converged again the day after (trial 153). The outliers, both with and without RC, can potentially be avoided by increasing the systems saturation bounds.

Fig. 9 shows the reference signal and the modified reference signal for two consecutive trials, indicating the continuous modifications to the reference signal for the inner PI control, which results in better tracking of the original reference $T_{z,r}$. The modified reference goes down during the morning hours when people arrive for work and increases after they leave in the afternoon. Furthermore, the lunch break causes a modification of the reference starting at around 11 am until 2 pm. This also correlates with the repeatable error shown in Fig. 6(a), when RC is not applied.

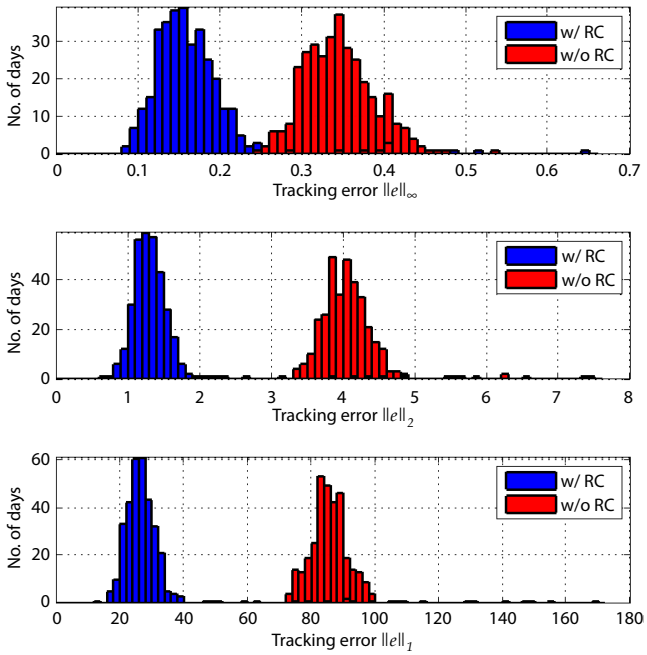

Fig. 7. Histogram of the ∞-norm, 2-norm, and 1-norm of the error for each day between 1st of Feb. to the 31st of Dec. with and without RC.

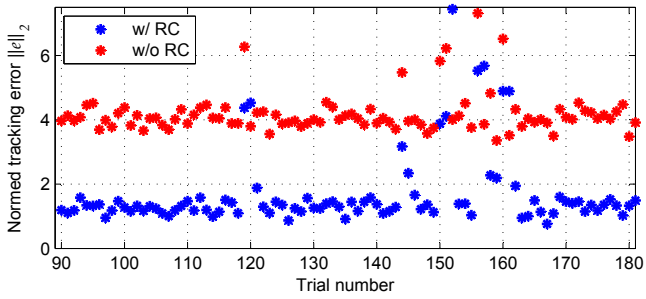

Fig. 8. 2-normed error for the days (trials) in May, June, and July with and without RC.

Fig. 10 shows the energy charge in U.S. dollars for each month of the year with and without RC using a time of use tariff effective in 2010 from the Arizona Public Service Company [18]. The expected energy consumption was estimated by assuming a constant heating efficiency of 1 and cooling efficiency of 2. The figure shows that there is negligible difference in energy charge with and without RC. Although there may be some uncertainty in the estimation of total consumed power since constant efficiencies are used, the importance is not in estimating the correct bill amount, but rather in comparing the control with and without RC under the same tariff schedule.

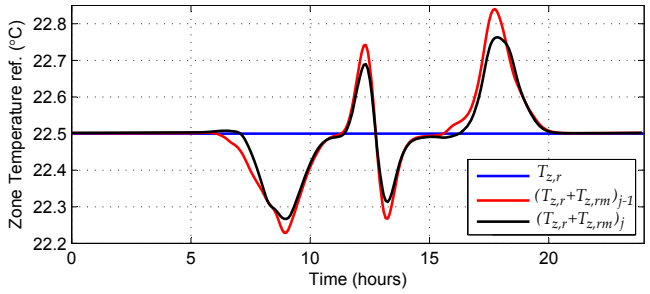

Fig. 9. Reference signal $T_{z,r}$ and modified reference signal for December 30th (red) and 31st (black) with $j = 334$ indicating the trial number.

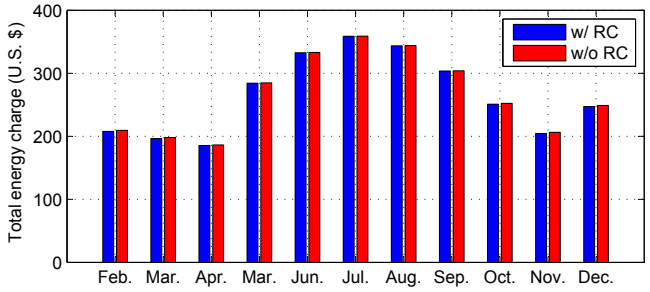

Fig. 10. Total energy charge each month with and without RC.

Table I summarizes the results and gives percentage improvements. We can conclude that the reference tracking

TABLE I

SUMMARY OF CASE STUDY RESULTS FOR SIMULATION FROM 1ST OF FEBRUARY TO THE 31ST OF DECEMBER.

quantity	without RC	with RC	% improvement
$\sum_{j=1}^{334} \|\mathbf{e}_j\|_\infty$	115.28	55.43	51.91
$\sum_{j=1}^{334} \|\mathbf{e}_j\|_2$	1357.70	475.87	64.95
$\sum_{j=1}^{334} \|\mathbf{e}_j\|_1$	29037.42	9584.77	66.99
Total energy (kWh)	50612	50450	0.32
Total cost (U.S. $)	2926.31	2915.54	0.37

performance was significantly improved with RC and this performance improvement came at no increase in energy consumption or spending cost. The tracking performance improvement could also be used to change the temperature reference set point to a value that requires less energy in terms of cooling or heating, while still staying within tolerable temperature bounds. This could be used in applications where we are mostly concerned with staying below a certain upper temperature bound, e.g. a food processing/storage facility or in a pharmaceutical processing facility. As a result, this could provide significant electricity savings.

VI. CONCLUSION

In this paper, the use of repetitive control is studied in the context of building temperature regulation. Improvements of more than 50% in set point tracking performance were achieved with no additional energy consumption, when the proposed RC procedure was applied to an existing feedback controller in a single zone office building simulation. No extra hardware is required and the proposed procedure is to use closed loop data from an appropriate time window, such as the previous year or month, and use that to calculate the repeatable-to-nonrepeatable ratio of the tracking error for each frequency. This is then used to design a repetitive controller with a learning filter that improves performance at frequencies where the repeatable part of the error is larger than the nonrepeatable part. The single zone case study used as an example is a simplified version of a true system, but detailed enough to give a realistic picture of the potential of learning control in building systems. Furthermore, the lumped parameter resistive-capacitive network modeling framework is applicable on larger buildings as well.

REFERENCES

[1] C. P. Underwood, *HVAC Control Systems: Modelling, Analysis and Design*. Spon Press, 1999.

[2] J. E. Seem, "A New Pattern Recognition Adaptive Controller with Application to HVAC Systems," *Automatica*, vol. 34, no. 8, pp. 969–982, August 1998.

[3] Q. Bi, W. Cai, Q. Wang, C. Hang, E. Lee, Y. Sun, K. Liu, Y. Zhang, and B. Zou, "Advanced controller auto-tuning and its application in HVAC systems," *Control Engineering Practice*, vol. 8, no. 6, pp. 633–644, June 2000.

[4] D. Lim, B. P. Rasmussen, and D. Swaroop, "Selecting PID Control Gains for Nonlinear HVAC&R Systems," *HVAC&R Research*, vol. 15, no. 6, pp. 991–1019, November 2009.

[5] F. Oldewurtel, A. Parisio, C. N. Jones, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, and M. Morari, "Use of model predictive control and weather forecasts for energy efficient building climate control," *Energy and Buildings*, vol. 45, pp. 15–27, February 2012.

[6] D. A. Bristow, M. Tharayil, and A. G. Alleyne, "A survey of iterative learning control," *IEEE Control Systems Magazine*, vol. 26, no. 3, pp. 96–114, June 2006.

[7] R. W. Longman, "Iterative learning control and repetitive control for engineering practice," *International Journal of Control*, vol. 73, no. 10, pp. 930–954, 2000.

[8] Y. Wang, F. Gao, and F. J. Doyle, "Survey on iterative learning control, repetitive control, and run-to-run control," *Journal of Process Control*, vol. 19, no. 10, pp. 1589–1600, December 2009.

[9] B. E. Helfrich, C. Lee, D. A. Bristow, X. H. Xiao, J. Dong, A. G. Alleyne, S. M. Salapaka, and P. M. Ferreira, "Combined H_∞-Feedback and Iterative Learning Control Design with Application to Nanopositioning Systems," in *American Control Conference*, Seattle, Washington, USA, June 2008, pp. 3983–3900.

[10] F. Lorenz and G. Masy, "Méthode d'évaluation de l'économie d'énergie apportée par l'intermittence de chauffage dans les bâtiments," *Traitement par differences finies d'un model a deux constantes de temps, Report No. GM820130-01. Faculte des Sciences Appliquees, University de Liege, Liege, Belgium*, 1982 (in French).

[11] M. Gouda, S. Danaher, and C. Underwood, "Building thermal model reduction using nonlinear constrained optimization," *Building and Environment*, vol. 37, no. 12, pp. 1255–1265, December 2002.

[12] G. Hudson and C. P. Underwood, "A simple building modelling procedure for MATLAB/SIMULINK," in *IBPSA Building Simulation Conference*, Kyoto, Japan, September 1999, pp. 776–783.

[13] National Renewable Energy Laboratory. Openstudio. http://openstudio.nrel.gov/.

[14] U.S. Department of Energy. Energyplus energy simulation software. http://apps1.eere.energy.gov/buildings/energyplus/.

[15] H. C. Hottel and A. F. Sarofim, *Radiative transfer*. McGraw-Hill New York, 1967.

[16] M. Tomizuka, T. Tsao, and K. Chew, "Discrete-Time Repetitive Controllers," *Journal of Dynamic Systems, Measurement, and Control*, vol. 111, pp. 353–358, September 1989.

[17] C. Kempf, W. Messner, M. Tomizuka, and R. Horowitz, "Comparison of Four Discrete-Time Repetitive Control Algorithms," *IEEE Control Systems*, vol. 13, no. 6, pp. 48–54, December 1993.

[18] APS, "Rate Schedule E-32TOU S - Small General Service (21 kW - 100 kW) - Time of Use," Arizona Public Service Company, Tech. Rep., 2010. [Online]. Available: http://www.aps.com/_files/rates/e-32TOUS.pdf