Peak Power Reduction for HVAC Operations in Multi-unit Commercial Buildings

Syed Ahsan Raza Naqvi*, Koushik Kar*, Saptarshi Bhattacharya† and Vikas Chandan†
*Rensselaer Polytechnic Institute, Troy, NY, USA, † Pacific Northwest National Laboratory, Richland, WA, USA.
Email: naqvis2@rpi.edu, koushik@ecse.rpi.edu, saptarshi.bhattacharya@pnnl.gov, vikas.chandan@pnnl.gov

***Abstract*—The load profiles of most commercial consumers are characterized by brief periods of very high power consumption followed by intervals of relatively lower demand. In order to flatten commercial load profiles, several power utilities in addition to billing energy consumption, levy a demand charge (DC) on the monthly peak demand. In this work, we consider the problem of joint optimization of energy costs (EC) and DC incurred by a multi-unit building which follows a demand response (DR) program. Despite the non-linear structure of the problem, we show how the optimal solutions can be obtained efficiently using linear programming. We evaluate the performance of the proposed power control scheme for various climate zones in the US. We show that depending on the ambient conditions and the prescribed tariff structure, our strategy can result in savings of up to nearly 19% compared to the baseline.**

I. Introduction

Power grids experience increased loads during summers as the demand for space cooling increases. In order to finance the maintenance of the infrastructure essential for delivering such high power, power utilities often levy a DC tariff on commercial consumers based on their monthly peak power demand [1]. Reports suggest that about 10–20% of the peak load in commercial buildings can be temporarily managed or curtailed to provide grid services [2]. Furthermore, effective building heating, ventilation and cooling (HVAC) controls can reduce power consumption by over 20% [3]. *Hydronic* (using water to provide heating/cooling) HVAC systems have become increasingly popular with commercial building operators due to their energy efficiency and potential for off-peak storage [4]. Through this paper, we aim to develop a rapidly deployable power control strategy for a grid-interactive efficient building (GEB) [5] that uses a hydronic HVAC system and is *DR-amenable*. We define DR-amenability as the practice of permitting indoor temperatures to vary within a temperature deadband. We aim to develop a power control strategy for commercial cooling that minimizes both the DC and the EC.

Peak load minimization in HVAC systems has received significant attention in the research community of late. For instance, in [6] the authors developed a grey-box model for characterizing a test-bed's thermal model. This model was used to determine the optimal pre-cooling required to achieve a desired temperature, while limiting the instantaneous power consumed to restrict the DC. In [7], the authors considered the problem of co-scheduling data-center and HVAC loads, while minimizing building EC and the DC. However, these papers do not consider the effects of heat transfer between adjacent zones in their system models. In contrast, our work tackles a more realistic scenario by taking into account this *thermal coupling* between spaces. This detailed heat transfer model can ensure that the simulated results of our control strategy are replicated with greater fidelity at the time of deployment. The authors in [8] formulated a bi-layer optimization framework for a variable air volume HVAC system to simultaneously minimize the thermal discomfort, EC and the DC. The work used a linear, discrete model to characterize the temperature in the building. Using time-sampled temperature values can introduce significant errors in computations. In order to address this issue, we develop a hybrid approach for estimating the temperature evolution in a building, that is neither purely discrete nor purely continuous.

In this work, we consider a hydronic HVAC system that serves a commercial, multi-unit GEB to meet its space cooling requirements. We develop an optimal control strategy that minimizes the EC and the DC incurred by this building. For our problem, we consider both the heat transfer between a unit and the outside environment, as well as that between adjacent units. The DR program is taken to allow the indoor temperatures to vary within a prescribed deadband, as agreed upon by the building operator and the power utility. Finally, we simulate our model to determine the optimal power required to achieve the desired cooling.

In the context of peak power reduction for thermostatically controlled loads (TCLs), our contributions are:

- We develop an optimization framework that jointly minimizes the EC and the DC incurred by a commercial GEB. The achievable peak power reduction indicates that a large-scale deployment of our control policy can potentially reduce the maintenance costs of the power delivery infrastructure, which may otherwise be adversely affected by large intermittent loads.
- Despite the apparent non-convex structure of the problem, we show how this framework can be expressed as a single vector of unknowns and subsequently solved using linear programming. This reduced dimensionality makes the approach scalable for large-scale deployment, and ensures operational convenience for the practitioner.
- We validate the performance of our framework for different climate zones as well as for various tariff structures.

II. Problem Formulation

In this work we consider a set of units $\mathcal{J}$ located in a commercial building. Each unit j has a pre-determined

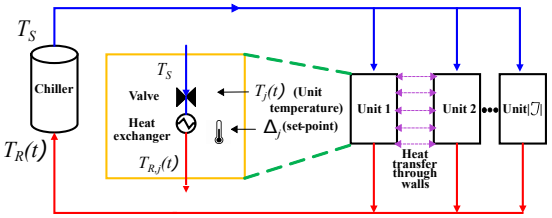

Fig. 1: Schematic of the hydronic HVAC system.

temperature set-point, Δ_j. Furthermore, ρ represents the maximum permitted temperature swing about Δ_j. This 'deadband' represents the flexibility in the HVAC load that is needed for implementing a DR program. The building is taken to be equipped with valves that can control the mass flow rate of the water depending on the cooling requirements of each individual unit. We assume that the temperature of each unit is spatially uniform, and use a typical R-C model to represent the thermal properties of the indoor units. The walls between adjacent rooms are taken to be thin and modeled as a single resistive element. It is further assumed that the these thermal parameters are known to (or well-estimated by) the utility. The HVAC system is assumed to operate at 100% efficiency, i.e., the electrical power used to cool the water gets fully utilized into thermal energy required to cool the units. The power consumed to cool unit j at time t is given by $P_j(t)$. The temperature of each unit is denoted by $T_j(t)$. Lastly, $P^{\text{total}}(t)$ is the total power required to cool all the units at time t.

We consider a hydronic HVAC system in which cooled water is distributed from a central source to spaces within a building. We take the central source to be a chiller powered by electricity. The temperature of the water leaving the chiller is denoted by T_S and is assumed to be a constant. The HVAC system is equipped with valves to control the mass flow rate of water at time t to the fan coil unit in space j, $\dot{m}_j(t)$. The temperature of the water returning to the chiller from unit j at time t is denoted by $T_{R,j}(t)$. Fig. 1 shows a simplified schematic of the HVAC system being studied here.

We pose our objective as an optimization problem that aims to minimize the weighted sum of (i) EC and (ii) the DC, over a pre-determined time horizon. We aim to use an MPC-based approach to determine the optimal power to achieve this objective. Mathematically, the objective is to minimize,

$$\alpha \int_0^\tau \pi(t) P^{\text{total}}(t) dt + (1-\alpha)\varpi[\max_{0\le t\le\tau} |P^{\text{total}}(t)|], \qquad (1)$$

subject to,

$$(\text{C}_1') \quad \dot{T}_j(t) = \frac{1}{C_j}\Big\{\frac{1}{R_j^o}[T_\infty(t) - T_j(t)] + P_j(t) + \sum_{i\in\mathcal{B}_j} \frac{1}{R_{i,j}^w}\Big(T_i(t) - T_j(t)\Big)\Big\}, \quad \forall j \in \mathcal{J},$$

$$(\text{C}_2') \quad P_j(t) = \dot{m}_j(t) S_p (T_S - T_{R,j}(t)), \quad \forall j \in \mathcal{J},$$

$$(\text{C}_3') \quad P_j(t) = h_R\Big(\frac{T_S + T_{R,j}(t)}{2} - T_j(t)\Big), \quad \forall j \in \mathcal{J},$$

$$(\text{C}_4') \quad 0 \le \dot{m}_j(t) \le \frac{1}{\phi_j}, \quad \forall j \in \mathcal{J},$$

$$(\text{C}_5') \quad -\rho \le T_j(t) - \Delta_j \le \rho, \quad \forall j \in \mathcal{J},$$

where α is the weighting parameter, ϕ_j is a constant and τ is the length of the prediction window. The initial temperature of unit j is $T_j(0) = T_j^0$. Additionally, $\pi(t)$ denotes the energy price at time t, while ϖ is the DC tariff. Let C_j represent the thermal capacitance of unit j. R_j^o and $R_{i,j}^w$ are the thermal resistances of the walls connecting unit j with the outside environment and with unit i, respectively. Let $\mathcal{B}_j$ denote the set of units that share a wall with unit j. Moreover, $T_\infty(t)$ is the ambient temperature at time t and is assumed to be known *a priori*. Let S_p be the specific heat capacity of water. Furthermore, T_S is the constant temperature of the water at the energy source, whereas $T_{R,j}(t)$ is the temperature of the water returning from the building at time t. Finally, h_R is the heat exchange coefficient between water and air.

For the objective function in (1), (C_1') models the temporal evolution of unit j's temperature. Constraint (C_2') links the power consumed with the mass flow rate and the change in temperature of the water. Constraint (C_3') expresses the power consumed in terms of the heat exchange coefficient and the temperature difference between the cooling water and the surrounding air in the unit. The temperature of the water is approximated to be the average of the temperature of the supplied and returning water. Constraint (C_4') enforces upper and lower bounds on $\dot{m}_j(t)$. Finally, constraint (C_5') limits the temperature swing about the set-point to be at most ρ° C.

The temperature of the water returning from the units to the chiller at time t is given by $T_R(t)$. This quantity can be obtained by the following equation [15]:

$$T_R(t) = \frac{\sum_{j\in\mathcal{J}} \dot{m}_j(t) T_{R,j}(t)}{\sum_{j\in\mathcal{J}} \dot{m}_j(t)} \qquad (2)$$

III. Analysis

Inspecting the constraints for the objective in (1), we notice that (C_2') makes the problem non-convex. Additionally, the problem contains several variables (like $\dot{m}_j(t)$, $P_j(t)$, $T_j(t)$ and $T_{R,j}(t)$) that must be determined optimally. Since we wish to study the efficacy of using commercial TCLs for DR, we attempt here to reformulate our problem and express it in terms of a single control knob, $P_j(t)$. We also express $T_j(t)$ in terms of known constants and $P_j(t)$. In the process, we express our problem as a convex problem with linear constraints. We first manipulate the constraint (C_2') to give,

$$\frac{P_j(t)}{\dot{m}_j(t) S_p} = T_S - T_{R,j}(t). \qquad (3)$$

Also, constraint (C_3') can be manipulated to give,

$$2\Big(\frac{P_j(t)}{h_R} + T_j(t)\Big) = T_S + T_{R,j}(t). \qquad (4)$$

Using (C_4') and setting $X_j(t) = \frac{1}{\dot{m}_j(t)}$, we get $\phi_j \le X_j(t) < \infty$. Adding (3) and (4) and using the lower bound of $X_j(t)$, we get,

$$P_j(t) \le G[T_S - T_j(t)], \qquad (5)$$

where $G = \frac{2S_p}{\phi_j + \frac{2S_p}{h_R}}$ is a constant.

Following the manipulations above, our problem is simplified to minimizing (1) subject to constraints (C_1'), (C_5') and (5). This results in two continuous variables, $T_j(t)$ (the state variable) and $P_j(t)$ (the control variable). However, in practice $\dot{m}_j(t)$, and hence $P_j(t)$, typically change only at discrete time intervals. Therefore, we can discretize the time scale for the valve operation into time instances. Each time instance is denoted by k and its duration is given by μ. We consider a total of K time instances where $K = \frac{\tau}{\mu}$. Thus, we express our control variable at time instance k as $P_j(k)$. Since T_∞ evolves slowly, the variation during a time instance is expected to be small. Hence, we represent it as a discrete variable $T_\infty(k)$.

The indoor temperature evolution is modeled as,

$$\dot{\mathbf{T}}(t) = \mathbf{A}\mathbf{T}(t) + \mathbf{H}(t), \tag{6}$$

where $\mathbf{T}(t) \in \mathbb{R}^{|\mathcal{J}|\times 1}$ column vector of temperatures of the unit at time t, $\mathbf{H}(t) \in \mathbb{R}^{|\mathcal{J}|\times 1}$ is a column vector given by $\{\frac{1}{C_j}(\frac{T_\infty(t)}{R_j^o} + P_j(t))\}$ and $\mathbf{A} \in \mathbb{R}^{|\mathcal{J}|\times|\mathcal{J}|}$ has elements $A_{i,j}$ defined as, $-\frac{1}{C_j}\left(\frac{1}{R_j^o} + \sum_{i\in\mathcal{B}_j}\frac{1}{R_{i,j}^w}\right)$ if $i = j$, and $\mathbb{1}_{i\in\mathcal{B}_j}\frac{1}{C_j}R_{i,j}^w$ otherwise. Here $\mathbb{1}_{i\in\mathcal{B}_j}$ is 1 if $i \in \mathcal{B}_j$ and is 0 otherwise.

Using time-sampled versions of $T_j(t)$ can introduce significant errors in our computations. Hence, we now express the continuous-time state variable in terms of $P_j(k)$. Taking $k = \lfloor \frac{t}{\mu} \rfloor$ and using [9], the building temperature at time t is,

$$\mathbf{T}(t) = \mathbf{M}(t-k\mu)\mathbf{M^{-1}}(0)\prod_{n=1}^{k-1}\Big(\mathbf{M}(\mu)\mathbf{M^{-1}}(0)\Big)\mathbf{T}_0 + \tag{7}$$

$$\sum_{k'=0}^{k-2}\Big(\prod_{n=1}^{k-k'-1}\mathbf{M}(\mu)\mathbf{M^{-1}}(0)\Big)\mathbf{I}\int_0^\mu \mathbf{M}(\mu)\mathbf{M^{-1}}(s)\mathbf{H}(k')ds +$$

$$\int_0^{t-k\mu}\mathbf{M}(t-k\mu)\mathbf{M^{-1}}(s)\mathbf{H}(k-1)ds,\ t\in[k\mu,(k+1)\mu],$$

where $\mathbf{I} \in \mathbb{R}^{|\mathcal{J}|\times|\mathcal{J}|}$ is the identity matrix, $\mathbf{T}_0 \in \mathbb{R}^{|\mathcal{J}|\times 1}$, with element T_j^0, denotes the temperature in all units at the beginning of the prediction horizon and $\mathbf{M}(t) \in \mathbb{R}^{|\mathcal{J}|\times|\mathcal{J}|}$ is the fundamental matrix for the system, $\dot{\mathbf{T}}(t) = \mathbf{A}\mathbf{T}(t)$. The matrix $\mathbf{M}$ is defined as:

$$\mathbf{M} = \Big[\lambda_1\mathbf{v}_1 \ldots \lambda_{|\mathcal{J}|}\mathbf{v}_{|\mathcal{J}|}\Big], \tag{8}$$

where λ_e and $\mathbf{v}_e$ are the e^{th} eigenvalue and eigenvector of matrix $\mathbf{A}$, respectively. Further details on the derivation of (7) are given in the detailed technical report for this work [16].

The optimization problem can now be expressed as,

$$\alpha\sum_{k=1}^{K}\{\mu\pi(k)P^{\text{total}}(k)\} + (1-\alpha)\varpi[\max_{0\le k\le K}|P^{\text{total}}(k)|], \tag{9}$$

subject to,

$$(C_1) \quad -\rho \le T_j(k\mu) - \Delta_j \le \rho, \qquad \forall j\in\mathcal{J}, k\in[0,K],$$
$$(C_2) \quad P_j(k) \le G[T_S - T_j(k\mu)], \qquad \forall j\in\mathcal{J}, k\in[0,K],$$
$$(C_3) \quad (7), \qquad k\in[0,K].$$

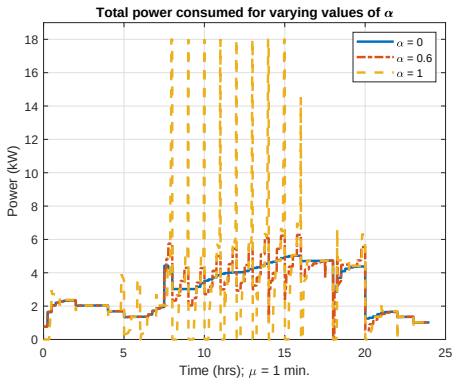

(a)

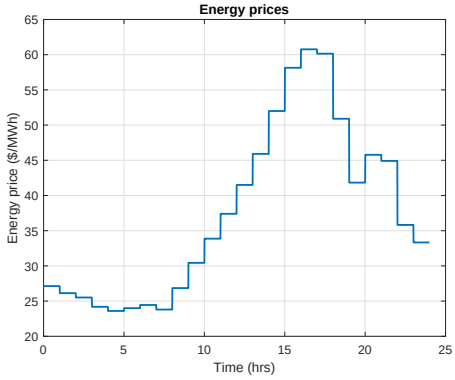

(b)

Fig. 2: (2a): Total power consumed for α=0, α=0.6 and α=1; (2b): Hourly energy pricing signal.

Henceforth, we will refer to (9) as the *energy costs plus peak demand charge minimization* (ECPDCM) problem. The formulation above has been expressed in terms of only $P_j(k)$. Furthermore, the cost function for the ECPDCM problem as well as the constraints are linear in $P_j(k)$.

IV. Numerical Study

In this section, we simulate our model to determine the power required to maintain unit temperatures within permitted bounds. Each unit is constrained to use a maximum power of 3 kW at any given time. The building follows a DR program whereby ρ = 2°C from 0000 hrs to 0759 hrs. As commercial buildings are usually occupied between 0800 hrs and 1800 hrs, we use ρ = 1°C for this interval. The baseline case for our simulations is when ρ = 0°C, i.e., when the building is not DR-amenable. The values of the remaining simulation parameters are recorded in [16].

A. Description of the Test-bed

We evaluate the performance of our power control strategy using a six-unit indoor space model, which represents a typical office-space, based on our testing facility located in Watervliet, NY. Refer to [10] for further details about the facility.

B. Performance Evaluation for Varying System Parameters

In this subsection, we use the hourly ambient temperatures observed in Albany, NY, on July 20th, 2019 [11]. We consider hourly energy prices as promulgated by NYISO for this day [12]. Here, ϖ is \$9.3/kW, which is the mean of the DC tariffs

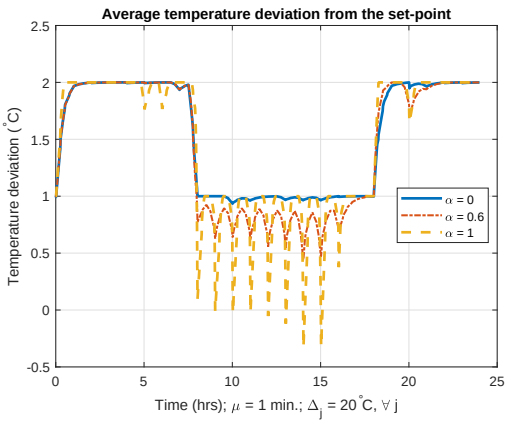

Fig. 3: Temperature deviation for α=0, α=0.6 and α=1, averaged for all six units in the test-bed.

charged by all utilities in NY [13]. We assume that the tariffs are known (or are well-estimated) in advance.

Fig. 2a shows how the total power consumed by the hydronic HVAC system operating under the ECPDCM strategy varies during the day, given the hourly pricing signals (Fig. 2b). The total power consumed by the building is sampled at 1-minute intervals. We use α=0 to represent the scenario where only the DC is minimized. When α=1, only the ECs are minimized. Finally, α=0.6 represents a compromise between EC and DC savings. It is noteworthy that between 0000 and 0500 hrs both the energy prices and the ambient temperatures progressively decrease. Hence, the power consumed exhibits a decreasing trend for all values of α under consideration. Later, an increase in hourly energy prices is preceded by instances of high power consumption of approximately 18 kW for $\alpha = 1$. These spikes in power consumption are followed by instances with zero power consumption. The units are pre-cooled prior to an increase in energy prices, thereby reducing power consumption during the intervals with higher energy prices, while also elevating the DC. The pre-cooling is less pronounced for α=0.6. Finally, no pre-cooling takes place for α=0 as this strategy is agnostic to changes in energy prices. The curve for α=0 shows a sharp increase just before 0800 hrs, signifying that the space is occupied and ρ is now 1°C.

Fig. 3 shows the temperature deviation from Δ_j (averaged over all units) when α=0, α=0.6 and α=1. It may be observed that for all three values of α, the temperature deviation remains close to ρ=1°C during periods of occupancy. This value is closer to ρ=2°C when the building is unoccupied. Thus, DR-amenability allows our control strategies to save energy without significantly compromising on the occupants' comfort. Furthermore, the approaches using α=0.6 and α=1 employ a pre-cooling mechanism to reduce power consumption during intervals with higher energy prices. This causes a reduction in the deviation from the set-point for the two control strategies prior to increases in energy prices. As previously explained, no pre-cooling occurs when α=0. Hence, this approach does not cause any notably sharp changes in temperature deviation.

Table I records how the system is affected by changing α for given $\pi(\cdot)$ and ϖ. Since the power meter records an average power flow for each 15-minute interval, the DC is based on the largest average 15-minute power flow during the billing period [14]. The product of this average power flow and ϖ gives the DC. We assume that the particular day being considered here sees the largest monthly load. The daily energy consumption is taken to be identical throughout the month. The monthly EC is the payable amount for the monthly energy consumption based on the per-unit energy prices. Lastly, the average temperature deviation represents the deviation of the mean temperature of all six units from Δ_j averaged over the entire 24-hour period.

Table I shows that increasing α decreases ECs, while also causing the pre-cooling to be more drastic. Therefore, the largest average 15-minute power flow, Λ, and hence the DC, increases in magnitude with increase in α. For the given tariff structure, the decrease in EC for increasing α does not compensate for the increase in DC. Since the DC dominates the total utility bill in this case, a general increase in the consumer's bill may be observed for increasing α. The absence of pre-cooling for the baseline causes (i) the DC to be significantly lower than those for α>0, and (ii) ECs to be the greatest as the temperature remains close to the set-point. We define % financial savings as the ratio of the difference between the total utility bills for the baseline and the ECPDCM strategy for a particular value of α, to the baseline utility bill. The results show that the % financial savings progressively decreases for α>0.2 due to a concomitant increase in the DC. Beyond α=0.6, ECPDCM does not result in any savings.

C. Performance Evaluation for Various Climate Zones

In this subsection, we determine the performance of the ECPDCM strategy for different climate zones and tariff structures. Here, we consider the hottest days of 2019 for three climate zones – 'Hot-Dry', 'Marine' and 'Very Cold'. For details on the representative tariff structures and ambient conditions for each zone, see [16]. While both Hot-Dry and Marine zones have on- and off-peak hours, the energy prices in the Very Cold zone are flat. We test our control strategy for α=0, α=0.2 and α=1 using representative values for $\pi(\cdot)$ and ϖ. We include the performance of α=0.2 as it resulted in the lowest utility bill in Table I. We wish to determine the effect of changing α on the energy consumption, peak demand and the utility bill for the three climate zones being studied here. These climate zones represent high (Hot-Dry), medium (Marine) and low (Very Cold) power demand for cooling.

Table II records the performance of our power control strategy for α=0, α=0.2 and α=1 for the three climate zones. The DC, the total utility bill and the % financial savings compared to the baseline have been determined in the same manner as in Section IV-B. As expected, the energy consumed for thermal regulation increases with increase in the average ambient temperature. It is noteworthy that although the energy consumption for α=0.2 in the Hot-Dry climate zone is higher than that for α=0, the former results in a lower monthly EC than the latter. This trend may be attributed to the pre-cooling action seen for α>0. It may also be seen that the ambient temperature and the tariff structure affect the peak power consumption for all three values of α. For instance, the values of Λ recorded for the Very Cold climate are the lowest among

TABLE I: Performance evaluation of the ECPDCM strategy for varying values of α.

	Baseline	$\alpha = 0$	$\alpha = 0.2$	$\alpha = 0.4$	$\alpha = 0.6$	$\alpha = 0.8$	$\alpha = 1$
I - Monthly energy consumed (kWh)	2318	1964	1964	1963	1962	1950	1894
II - Monthly energy cost ($)	94.2	81.0	81.0	80.7	80.2	79.4	76.9
III - Peak 15-min. power flow (kW)	5.36	5.02	5.02	5.87	6.17	8.36	9.78
IV - Demand charge ($) [**III** $\times \varpi$]	49.9	46.7	46.7	54.6	57.4	77.8	90.9
V - Total utility bill [**II** + **IV**] ($)	144.1	127.7	127.6	135.5	137.6	157.2	167.8
VI - % financial savings	-	11.38	11.45	5.97	4.51	-9.09	-16.45
VII - Avg. temperature deviation (°C)	0	1.55	1.53	1.50	1.45	1.44	1.47

TABLE II: Performance evaluation of the ECPDCM strategy for α=0, α=0.2 and α=1 for various climate zones and tariff structures.

	Hot-Dry			Marine			Very Cold		
	α=0	α=0.2	α=1	α=0	α=0.2	α=1	α=0	α=0.2	α=1
Monthly energy consumed (kWh)	2391	2397	2375	1668	1670	1669	621	618	616
Monthly energy cost ($)	470.80	468.90	465.30	86.48	85.74	85.74	36.77	36.57	36.44
Peak 15-min. power flow (kW)	6.75	10.06	11.42	5.73	9.82	11.44	3.02	3.02	3.05
Demand charge ($)	77.29	115.19	130.76	14.97	25.64	29.87	41.89	41.89	42.20
Total utility bill ($)	548.09	584.09	596.06	101.45	111.38	115.61	78.66	78.46	78.64
% financial savings	9.40	3.45	1.47	12.10	3.50	-0.16	18.59	18.80	18.61
Avg. temperature deviation (°C)	1.47	1.45	1.47	1.11	0.96	0.99	0.80	0.82	0.82

the three climate zones being studied. Furthermore, since this climate zone uses a flat pricing scheme, the power control scheme with $\alpha = 1$ does not pre-cool the units. Hence, there is no appreciable difference in Λ for the three values of α for this climate zone. In fact, it can be inferred that the DC incurred for the three values of α is not drastically different from that incurred for the baseline. Hence, unlike the other climate zones, the disparity in the monthly utility bills between the baseline and the system using ECPDCM in the Very Cold zone can be attributed to the difference between the ECs in the two cases. Therefore, the DR-amenability of our system results in significant savings in the Very Cold climate zone as compared to the other two zones. Moreover, % financial savings with respect to the baseline case for for Hot-Dry and Marine climate zone tends to decrease with increasing value of α. Finally, the average temperature deviation is observed to be lower for cooler climates. This is because the temperature gradient between the ambient and the units is less steep for climate zones with lower average ambient temperatures. Therefore, less heat is transferred from the surroundings to the units, thereby lowering the power demand for space cooling.

V. Conclusion

In this work, we aimed to determine the optimal power consumption of a commercial, multi-unit GEB that minimizes the weighted sum of EC and DC. We showed how this problem can be transformed into a convex optimization problem with linear constraints. We also tested our control strategy under high, medium and low cooling demand scenarios for various tariff structures. Our initial results showed that the proposed approach can potentially result in significant financial savings for DR-amenable commercial consumers. Our control policy can also be used in conjunction with various grid services, e.g. frequency regulation, at different temporal granularities. This has been left as future work.

References

[1] "Understanding demand charges," [Online] Available: nyserda.ny.gov/All-Programs/Programs/Energy-Storage/Energy-Storage-for-Your-Business/Understanding-Demand-Charges

[2] M. Piette, D. Watson, N. Motegi and S. Kiliccote, "Automated critical peak pricing field tests: 2006 pilot program description and results," *LBNL-59351*, 2007.

[3] N. Fernandez, S. Katipamula, W. Wang, Y. Huang and G. Liu, "Energy savings modeling of re-tuning energy conservation measures in large office buildings," in *Journal of Building Performance Simulation*, vol. 8, no. 6, pp. 391-407, 2014.

[4] U.S. Department of Energy, "Ductless hydronic distribution systems," Nov. 2011, [Online] Available: energy.gov/sites/prod/files/2013/12/f5/arbi_hydronic_webinar.pdf

[5] U.S. Department of Energy, "Grid-interactive efficient buildings," Apr. 2019, [Online] Available: energy.gov/sites/prod/files/2019/04/f61/bto-geb_overview-4.15.19.pdf

[6] A. Vishwanath , V. Chandan, and K. Saurav, "An IoT-based data driven precooling solution for electricity cost savings in commercial buildings," in *IEEE Internet Things J*, vol. 6, no. 5, Oct. 2019.

[7] T. Wei, M. Islam, S. Ren and Q. Zhu, "Co-scheduling of datacenter and HVAC loads in mixed-use buildings," in *Proc. International Green and Sustainable Computing Conference (IGSC)*, Nov. 2016, pp. 1-8.

[8] J. Mei and X. Xia, "Multi-zone building temperature control and energy efficiency using autonomous hierarchical control strategy," in *Proc. IEEE International Conference on Control and Automation (ICCA)*, Jun. 2018, pp. 884-889.

[9] "Nonhomogeneous linear systems of differential equations with constant coefficients," [Online] Available: people.math.gatech.edu/ xchen/teach/ode/NonhomoSys.pdf.

[10] M. Minakais *et al.*, "Design and instrumentation of an intelligent building testbed," in *Proc. IEEE Int. Conf. Automat. Sci. Eng. (CASE)*, Gothenburg, Sweden, Aug. 2015, pp. 1-6.

[11] "Weather in July 2019 in Albany, New York, USA," [Online] Available: timeanddate.com/weather/usa/albany-ny/historic?month=7&year=2019

[12] "Day-Ahead Market LBMP - Zonal", [Online] Available: mis.nyiso.com/public/P-2Alist.htm

[13] National Renewable Energy Laboratory (NREL), "A Survey of U.S. demand charges," Sep. 2017.

[14] Northwestern Energy, "Demand charges explained," [Online] Available: northwesternenergy.com/docs/default-source/documents/E-Programs/E-demandcharges.pdf

[15] S. Bhattacharya, V. Chandan, V. Arya, K. Kar, "DReAM: Demand response architecture for multi-level district heating and cooling networks," *e-Energy*, May 2017, pp. 353-359.

[16] www.dropbox.com/s/u66y0obn06aho1z/ISGT_2021.pdf?dl=0