Comparative Evaluation of Model Predictive Control Strategies for a Building HVAC System

Vamsi Putta[1], Guangwei Zhu[1], Donghun Kim[2], Jianghai Hu[1], and James Braun[2]

Abstract— **The paper presents an evaluation procedure for a few promising Model Predictive Control (MPC) based approaches in supervisory control of building heating, ventilation and air-conditioning (HVAC) systems. A case study is established and control trajectories are generated using the proposed solvers. The resulting trajectories are evaluated in terms of benchmark metrics taking energy costs and occupant discomfort into account. A direct approach to incorporate occupant discomfort in the cost function is also proposed. Results indicate potential savings using MPC solvers when compared to the baseline.**

I. Introduction

Model predictive control (MPC) is increasingly being viewed as a practical solution for building heating, ventilation and air-conditioning (HVAC) systems control [1], [2]. The ability to incorporate information such as weather forecasts and occupancy profiles in real time decision making makes MPC approaches highly attractive in this regard. However, the complexity of building models can make such approaches infeasible for all but the smallest buildings. Hence, some simplifying assumptions are usually made when formulating the MPC problem for building controls.

The different assumptions and simplifying techniques lead to different control strategies or solvers. If the underlying assumptions are too different, comparison of the performance of the solvers is no longer straightforward. Hence there is a need for a common benchmark applicable to all solvers that can rank solvers based on their "optimality". Such a benchmark would be ideally be a critical part of a tool chain for designing implementable optimal control solutions in buildings. Additionally, such a benchmark can also help evaluating the benefits of retrofitting buildings.

In this paper, we establish a simple case study to study the performance of a few promising model predictive control strategies. Benchmark metrics corresponding to energy costs and losses due to thermal discomfort are also formulated. The benchmark metrics themselves are utilized in formulating the optimization problems for the MPC based solvers. In particular, we propose an approach to quantify losses due to thermal discomfort that can be incorporated into cost functions. Using the benchmark metrics, the performance of the solvers are quantified and compared. A conventional control profile is used as the baseline from which the savings gained by the use of MPC are evaluated. Computational complexity and implementation issues of the solvers are also discussed.

The paper is organized as follows. In Section II, we discuss the building and HVAC system models and the assumptions used for the case study. We also formulate the cost functions used to benchmark the different strategies. A baseline conventional controller is discussed and simulation results presented in Section III. We describe different MPC strategies in Section IV. We analyze and remark on the solver performances in Section V. Finally, we make some concluding remarks in Section VI.

II. Case Study

A. Modeling

For the purpose of this study, a simplified state-space model of the Purdue Living Lab (facility under construction) was used as the common test case. The model was obtained by applying energy balances at suitably chosen points inside the room and the walls. The non-linear model thus obtained captured all the transients due to heat extraction and solar radiation. Details of the modeling procedure are reported in [3]. For application in model predictive approaches, the non-linear model was simplified. After linearization, discretization and appropriate model order reduction techniques, the model may be expressed in the form [3]

$$\begin{aligned} x_{(t+1)} &= Ax_t + Bu_t + Fw_t \\ y_t &= Cx_t, \end{aligned} \tag{1}$$

where A, B, F and C represent the system matrices of reduced dimension obtained via model order reduction and t denotes the discrete time instant. The state vector x_t represents a transformed vector containing information about the temperatures of the wall and air nodes. The controllable inputs (rate of heat extraction kW) are denoted by the vector u_t. Vector w_t denotes the exogenous (uncontrollable) inputs acting on the envelope (solar radiation, ground radiation). The output matrix C provides the linear relation between the states and the outputs. For the Living Lab model, two outputs namely the zone air temperature T_z and mean radiant temperature T_{rad} were considered available. The linear model obtained after model order reduction was validated against the high fidelity non-linear model with negligible error [3].

A quasi-static approach was used in modeling the Air Handling Unit (AHU) and cooling plant supplying cool air to the room. This was motivated by the need to reduce the

[1]Vamsi Putta, Guangwei Zhu and Jianghai Hu are with the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47906, USA. {vputta,guangwei,jianghai} at purdue.edu

[2]Donghun Kim and James Braun are with the School of Mechanical Engineering,Purdue University, West Lafayette, IN 47906, USA. {kim1077,jbraun} at purdue.edu

large number of controllable variables available on the plant side. Under certain simplifying assumptions such as constant efficiency and pressure ratios, the power consumption of the different components (namely coils, fans and pumps) were modeled using empirical models. The total power consumption of the plant was generated by summing up the individual consumptions and had the following form

$$P = f(\dot{m}_{vent}, T_{vent}, T_z, T_{OA}, RH). \quad (2)$$

Here, $\dot{m}_{vent}$ and T_{vent} denote the flow rate and the temperature of the air supplied by the AHU while T_z and T_{OA} denote the room air temperature and the ambient temperature respectively. RH stands for the relative humidity of the ambient air. Noting that the controllable variables $\dot{m}_{vent}$ and T_{vent} together affect the room dynamics in (1) via the rate of heat extraction u, we can relate the AHU operating power to the heat extraction rate using a lookup table of the form

$$P^*(u; T_z, T_{OA}, RH) = \min_{u=g(\dot{m}_{vent}, T_{vent})} P(\dot{m}_{vent}, T_z, T_{vent}, T_{OA}, RH) \quad (3)$$

The lookup table $P^*(u; T, T_{OA}, RH)$ gives the minimum power consumption of the AHU when supplying a heat extraction rate of u at zone temperature T, ambient temperature T_{OA} and RH humidity level. Using this optimal map reduces the degrees of freedom ($\{\dot{m}_{vent}, T_{vent}\}$ to u) available in controlling the room dynamics. This leads to a reduction in the search space when searching for optimal plant operation. Also, the lookup table in effect decouples the optimal control of the room dynamics and the problem of choosing the plant set points. It should be noted that there is an implicit assumption that the plant components have much smaller time constants compared to the zone dynamics to justify the usage of a quasi static model. Further details on the plant modeling are given in [4].

In the following section, we motivate the necessity of a common metric in evaluating different MPC solvers and formulate one such metric.

B. Benchmark Metric

Optimal HVAC operation involves the twin objectives of maintaining occupant comfort while minimizing the energy costs. As the objectives are typically conflicting, the optimization requires finding a suitable trade off. Despite the different formulation of MPC solvers, they can be rated in terms of their performance with respect to the two objectives. Additionally, the computational costs of each solver must also be taken into account.

Computing the energy costs incurred due to an optimal trajectory is relatively straight forward. From the lookup table P^* described in section II, it is possible to compute the power costs corresponding to a control trajectory u as follows

$$J_{power} = \sum_t r_t P^*(u; T_z, T_{OA}, RH)_t. \quad (4)$$

Here r_t denotes the time-of-day price in dollars per unit power per time step and $P^*(\cdot)_t$ refers to the computed power consumption at time t. We assume that the complete information of the ambient conditions are known and the zone temperature T_z is available from the room dynamics as component of the output y_t. Using (4) it is possible to compute the actual energy costs incurred by the trajectory generated by any particular solver. We ignore the peak demand costs in this study, to simplify computation.

The choice of occupant discomfort metric is more involved due to the hidden nature of costs incurred due to thermal discomfort. Thermal discomfort is reflected in the decreased productivity of the occupants, resulting in indirect losses. Hence a reasonable performance metric of any control strategy would be the losses incurred due to productivity loss. Such losses can be quantified as the

$$J_{discomfort} = \sum_t S_t \cdot \eta_t, \quad (5)$$

where S_t represents the average salary of the occupants in the buildings at time t and η_t denotes the productivity loss due to thermal discomfort in percent at time t. During unoccupied periods the average salary S_t is set to 0 which avoids prioritizing comfort unnecessarily. Productivity losses are typically obtained through regression models once thermal discomfort is quantified [5]. Predicted Mean Vote (PMV) and the related Predicted Percentage Dissatisfied (PPD) are two models that calculate discomfort in terms of the prevailing conditions accurately [6]. The PMV model uses information about the room air temperature T_z, mean radiant temperature T_{rad} along with other factors to quantify comfort on a continuous scale ranging from -3(very cold) to 3(very hot). A PMV value of zero corresponds to the best average thermal comfort among occupants. The PPD scale predicts the percentage of people dissatisfied with the thermal conditions and is related to the PMV as follows [6]

$$PPD = 100 - 95e^{-(0.03353\,PMV^4 + 0.2179\,PMV^2)}. \quad (6)$$

The PPD equation predicts that at least 5% of the occupants will be dissatisfied under any thermal conditions. Productivity losses in employees was as derived as a function of the PMV (PPD) in [5]. The variation of productivity loss η with PMV are depicted in Fig. 1. With PMV computed from the output trajectory y_t, it is now possible to assign a dollar cost to the losses due to discomfort. It must be observed that the choice of the discomfort metric now allows a direct comparison with the power cost due to the same units (dollars). This obviates the need for scaling factors when comparing the magnitude of savings in both metrics and lends more insight into the trade off between the conflicting objectives.

Traditional controllers are designed to track zone air temperature set points. As the productivity loss (PMV/PPD) depends upon both the zone air temperature and mean radiant temperature, set point tracking might be suboptimal with respect to occupant productivity. For the same reason, using a PMV/PPD based cost metric in optimization can lead to more savings by utilizing the energy storage inside the building walls (reflected in the mean radiant temperature).

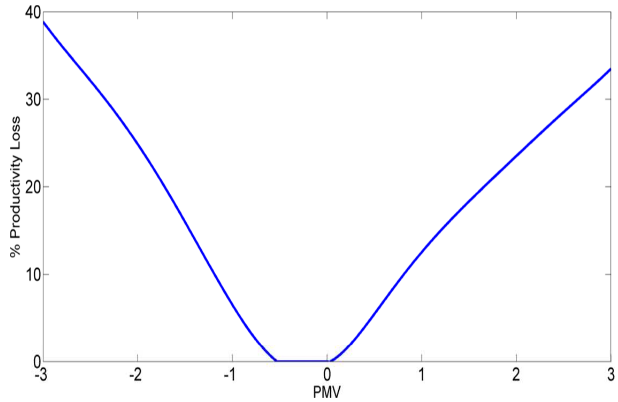

Fig. 1: η (Productivity Loss %) vs. PMV

Additional factors influencing the choice of MPC strategies are the computational costs involved and ease of implementation. Since computational costs vary with implementation, a direct comparison need not be accurate. However, the computational time taken is a good measure of the computational complexity, especially when run on similar computing platforms.

The next section summarizes the parameters used for the case study. All the control strategies reported were tested for the same parameters facilitating direct comparison.

C. Case Study- Other Factors

A 12 dimensional state space model of the Purdue Living Lab model was used for the room dynamics. The sole controllable input u was the rate of heat addition/extraction from the zone. The effects of the ambient environment was captured in 16 exogenous variables w whose complete data was assumed to be known. For simulations, this data was extracted from Indiana TMY2 weather data (available for every 10 min intervals) corresponding to July 2010. The state space model was discretized with a time step of 10 minutes assuming zero order hold. The choice of the time step was motivated by the resolution of the weather data available. The building was assumed to be occupied daily from 7 am to 6 pm by 20 occupants. The average salary of the occupants was chosen to be $70,000$ \$/year (or $S_t = 1.92$ \$/10 minutes). Utility price variation was considered with a peak price of 0.03 \$/kWh ($r_t = 0.005$ \$/kW/10 minutes) daily from 10 am to 3 pm. Nominal price during off peak hours was fixed at 0.01 \$/kWh.

PMV (PPD) computation requires additional information about occupant activity and room conditions apart from zone and mean radiant temperature. These were assumed to be constant for all the simulations. The occupant metabolic activity, clothing level was fixed at 1 met and 0.5 clo which are typical values for non-vigorous activity in summer. Air flow velocity was assumed to be constant (0.2 m/s) for all occupants. The ratio of surface areas of the clothed to the nude body was fixed at the common value of 1.1. Relative humidity inside the room was assumed to be the same as the ambient which is justified due to the little variation in ambient humidity during occupied periods.

Due to the choice of a summer weather profile, the AHU was constrained to operate in a cooling mode ($u \leq 0$). The power consumption of the AHU was neglected when no heat was being extracted ($u = 0$). This assumption was not too limiting as the coil power consumption dominated that of the fans in the plant model. Also, peak demand charges are ignored while calculating energy costs.

All simulations were performed on Intel Core2Duo desktops using MATLAB. The reported computation times are representative of the computational complexity of the solvers.

In the next section, we implement a conventional (non MPC) control law for the Living Lab case study to obtain a baseline performance metric. This will later be used to compare and contrast MPC based solver.

III. Conventional Control

For the Living Lab model described in section II, a setback strategy with proportional tracking was simulated. The control involved turning off the AHU during unoccupied periods and set-point tracking using a proportional controller during occupied controllers. To allow for precooling, the AHU was turned on an hour before the occupancy started. Time-of-day prices were not taken into account. The results of using a 26^oC set-point during occupancy are shown in Fig. 2. The results indicate there is no necessity of precooling

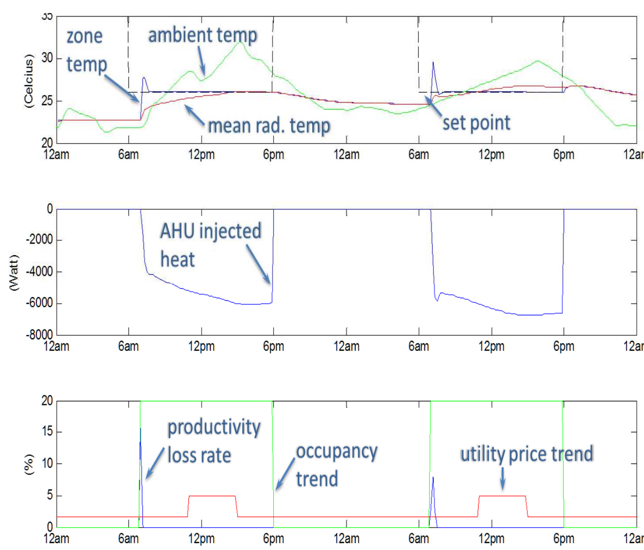

Fig. 2: Trajectories under conventional control.

before occupancy as the room temperature when allowed to float is lower than the setpoint of 26^oC. Evaluating the power and discomfort metrics in (4) and (5) respectively for the resulting state and control trajectories yields a power cost of $J_{power} = 0.287$ \$/day and a discomfort penalty of $J_{discomfort} = 0.720$ \$/occupant/day. These results quantify the baseline metrics which can be improved in some aspects by using MPC based approaches. The conventional control strategy has negligible computational complexity as there is no optimization involved at any time.

IV. Model Predictive Control Strategies

Model predictive control involves forecasting the system trajectories, over a prediction horizon L, and making an optimal decision based on the prediction. The first step of the optimal decision is applied to the system and the process repeated with an updated forecast. Assuming complete information of exogenous inputs w in (1), the predicted trajectories at time t under a control input trajectory $u_{t|t}, u_{t+1|t}, \ldots, u_{t+L-1|t}$ are described by

$$\begin{aligned} x_{t+k+1|t} &= Ax_{t+k|t} + Bu_{t+k|t} + Fw_{t+k|t} \\ T_{t+k|t} &= Cx_{t+k|t}, \qquad x_{t|t} = x_t, \quad k = 0, 1, \ldots, L \end{aligned} \tag{7}$$

where the subscript $k+t|t$ is used to denote the predicted value at time $k+t$ formed by propagating the initial value at time t. The model predictive control problem can be formulated using the predicted trajectories to define a optimization problem over the look ahead horizon L as follows.

$$u^* = \arg\min \sum_{k=0}^{L-1} J_k(x_{t+k|t}, u_{k+t|t}) + J_L(x_{t+L|t}) \tag{8}$$

J_k represents the cost incurred at k steps into the future. The optimization problem is constrained by the dynamics in (7). Out of the resulting optimal input sequence $u^*_{t|t}, u^*_{t+1|t}, \ldots, u^*_{t+L-1|t}$, only the first input $u^*_{t|t} = u_t$ is applied to the system in (1) and the trajectories are predicted with x_{t+1} as the initial condition. A major challenge in applying model predictive approaches to buildings is the computational complexity presented by optimization at every time step. As the cost functions J_i(typically representing power and thermal comfort) need not be convex, one needs to resort to numerical optimization. Additionally, the dimension of the search space grows linearly with the prediction horizon L as the length of the trajectory to be optimized grows. Hence, some approximations are usually made to facilitate real-time computation. We describe three different strategies (solvers) based on the approximations made and compare the performance in terms of the benchmark cost function when applied to the Living Lab case study.

A. *Affine Quadratic Regulator Based Solver*

Model predictive control based approaches are, in general, computationally expensive due to the necessity for numerical optimization. However, if quadratic approximations to the cost functions are available, a closed form solution to the optimization problem is possible, eliminating the need for numerical optimization routines. More specifically the cost function J_k in (8) must be a sum of quadratic functions of the following form

$$\begin{aligned} J_k &= (x_{t+k|t} - x^{\text{ref}}_{t+k})^T Q_{t+k}(x_{t+k|t} - x^{\text{ref}}_{t+k}) + \\ &\quad (u_{t+k|t} - u^{\text{ref}}_{t+k})^T R_{t+k|t}(u_{t+k} - u^{\text{ref}}_{t+k}), \\ &\quad k = 0, 1, 2, \ldots, L-1 \\ J_L &= (x_{t+L|t} - x^{\text{ref}}_{t+L})^T Q_L (x_{t+L|t} - x^{\text{ref}}_{t+L}). \\ &\quad \text{subject to (7).} \end{aligned} \tag{9}$$

Here Q_{k+t} (R_{k+t}) are positive semidefinite (positive definite) matrices for $k = 0, 1, \ldots, L-1 (k = 0, 1, \ldots, L)$ respectively. Intuitively, the quadratic term in the state and control inputs represent the predicted occupant discomfort and power consumption at time $t+k$. x^{ref} models a trajectory that meets the comfort requirement while u^{ref} can be used to track efficient control trajectories. The problem of minimizing a quadratic cost function J_k subject to affine dynamics of (7) is termed an Affine Quadratic Regulator (AQR) problem. To solve the problem, we invoke the principle of dynamic programming and assert that that for a particular sequence $u_{t|t}, u_{t+1|t}, \ldots, u_{t+L-1|t}$ to be optimal over the prediction horizon, any subsequence starting from a middle point u_{t+k} till the end must also be optimal for the corresponding shortened horizon subproblem. This principle can be expressed by the Bellman equation as,

$$\begin{aligned} &V_{t+k|t}(x_{t+k|t}) = \\ &\quad (x_{t+k|t} - x^{\text{ref}}_{t+k})^T Q_{t+k}(x_{t+k|t} - x^{\text{ref}}_{t+k}) + \\ &\quad \min_{u_{t+k|t}} \big\{ (u_{t+k|t} - u^{\text{ref}}_{t+k})^T R_{t+k}(u_{t+k|t} - u^{\text{ref}}_{t+k}) + \\ &\qquad V_{t+k+1|t} x(t+k+1|t) \big\}. \end{aligned} \tag{10}$$

The Bellman equation holds for all $k = 0, 1, \ldots, L-1$. The function $V_{t+k|t}(\cdot)$ is called the value function and has a general quadratic form

$$\begin{aligned} V_{t+k|t}(x_{t+k|t}) = x^T_{t+k|t} H_{t+k|t} x_{t+k|t} + g^T_{t+k|t} x_{t+k|t} + c_{t+k|t}, \\ k = 0, 1 \ldots, L. \end{aligned} \tag{11}$$

Noting that the constant $c_{t+k|t}$ does not affect the minimization in (10), and substituting for $V_{t+k+1}(x_{t+k+1|t})$ the quadratic form in (11), the optimization problem can be converted into optimizing a quadratic form. Minimizing the quadratic expression yields the following iterative expressions for $k = 0, 1, \ldots, L-1$.

$$\begin{aligned} H_{t+k|t} &= Q_{t+k} + M^T_{t+k|t} R_{t+k} M_{t+k|t} + S^T_{t+k|t} H_{t+k+1|t} S_{t+k|t}, \\ g_{t+k|t} &= -Q_{t+k} x_{t+k} + M^T_{t+k|t} R_{t+k}(N_{t+k|t} - u^{\text{ref}}_{t+k}) + \\ &\quad S^T_{t+k}(H_{t+k+1|t}[Fw_{t+k|t} + BN_{t+k|t}] + g_{t+k+1|t}), \\ u^*_{t+k|t} &= M_{t+k|t} x_{t+k|t} + N_{t+k|t}, \end{aligned} \tag{12}$$

where

$$\begin{aligned} P_{t+k|t} &= R_{t+k} + B^T H_{t+k+1|t} B, \\ M_{t+k|t} &= -P^{-1}_{t+k|t} B^T H_{t+k+1|t} A, \\ S_{t+k|t} &= A + BM_{t+k|t}, \\ N_{t+k|t} &= P^{-1}_{t+k|t}(R_{t+k} u^{\text{ref}}_{t+k} - B^T g_{t+k+1|t} - \\ &\quad B^T H_{t+k+1|t} F w_{t+k|t}), \end{aligned}$$

Iterating (12) backward starting with $H_{t+L|t} = Q_{t+L}$ and $g_{t+L|t} = 0$ it is possible to compute the optimal input sequence over the prediction horizon L.

In order to formulate the HVAC control problem in the AQR framework, it is imperative that quadratic approximations to the power cost and discomfort penalty be obtained. Using the lookup table P^* and the functional form of productivity loss, it is possible to obtain a family of positive definite quadratics to approximate the power consumption as

$$P^*(u; T_z, T_{OA}, RH) \approx (u - u^{\text{ref}})^T R(u - u^{\text{ref}}), \forall u.$$

Computing the matrices R that best fit the lookup table for every value of T_z, T_{OA}, RH yields the required family of matrices. At any time t, R_t in (9) is chosen from the family of quadratics based on the current values of T_z, T_{OA}, RH. This yields a close approximation of the actual power consumption of the AHU at time t.

The discomfort penalty $J_{discomfort}$ is expressed as a quadratic in a similar manner. Here the governing equation becomes

$$\eta(T_z, T_{rad}) \approx (x - x^{\text{ref}})^T Q(x - x^{\text{ref}}), \quad \forall T_z, T_{rad}, Cx = [T_z, T_{rad}]^T.$$

Again Q is parametrized in terms of any time varying exogenous factors (RH) and substituted in (9) based on the prevailing values. Fig. 3 show the validity of the quadratic fits for certain fixed parameters. The AQR formulation does not

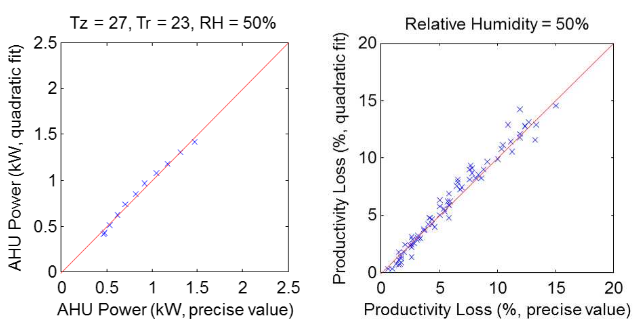

Fig. 3: Validation of quadratic fits

allow for non-linear constraints during optimization. Hence, the resulting control law is not guaranteed to yield $u_t \leq 0$ all the time. Thus, the AQR based solver might recommend switching between the AHU heating and cooling modes. However, the changes in the plant dynamics between these two modes cannot be adequately captured using a single family of quadratics. In the present case study, we use the heuristic of setting $u = 0$ when the AQR solver recommends a positive value of u. This is equivalent to saturation in the controller and results in suboptimal performance. However the simulation results, presented in Fig. 4, show little performance degradation. In fact, the benchmark metrics, calculated to be $J_{power} = 0.284$ \$/day and $J_{discomfort} = 0.517$ \$/day/occupant, indicate superior performance compared to conventional control. The existence of closed form solution implies high computational efficiency, making AQR feasible for real time implementation.

B. Sequential Quadratic Programming Based Solver

To formulate the MPC problem as an AQR problem requires a quadratic cost function. Sequential Quadratic Programming (SQP) relaxes the assumption by requiring

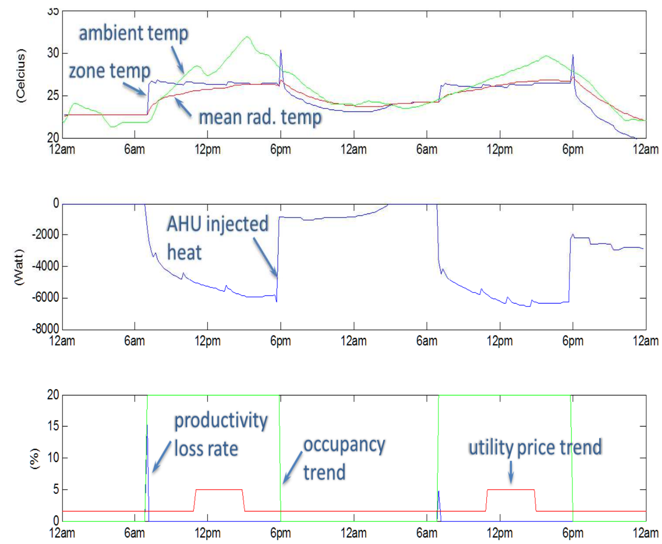

Fig. 4: Trajectories under affine quadratic regulator based solver.

the cost function to be only twice continuously differentiable. SQP iteratively approximates the cost function as a quadratic and the search direction is chosen to be the corresponding minimizer. SQP involves multiple numerical computations of the cost function Hessian. This can be prohibitively expensive when the cost function evaluation is time-consuming. This is especially true when computing the PMV or the PPD, which involve solving a set of non-linear algebraic equations at each step. To facilitate computing, a cubic regression model was developed for expressing the PPD in terms of the zone air temperature and the mean radiant temperature. Similarly, the lookup table P^* was also represented using a global quadratic $P^*(u; T_z, T_{OA}, RH) = [u, T_z, T_{OA}, RH]^T \mathbf{P}[u, T_z, T_{OA}, RH] + \mathbf{P}_0$. The MPC problem at time t was formulated with the cost function as

$$\min \sum_{k=0}^{L-1} J_k \quad \text{subject to}$$

$$PPD_{t+k} \leq \begin{cases} 10\% & \text{if occupied at time } t+k. \\ 30\% & \text{otherwise,} \end{cases}$$

where J_k is the power consumption at time $t + k$ computed from the global quadratic expression for P^*. The constraint on the PPD prioritizes comfort during occupied periods while allowing for less conservative cooling requirements during other times. The PPD threshold value (10%) was chosen close to a PMV value of 0.5 in order to minimize productivity loss during occupancy. The MPC problem is numerically solved using the SQP solver with a cubic regression model being used for PPD computation. Further details of the SQP solver are presented in [4]. Results of the simulation of the SQP solver for the case study are presented in Fig. 5. As in the case of conventional controller there is no significant precooling observed. However, the zone temperature at the start of occupancy is too high for occupant comfort causing a temporary productivity loss. The benchmark metrics were calculated to be $J_{power} = 0.274$ \$/day and $J_{discomfort} = 1.11$ \$/occupant/day. A prediction horizon of $L = 6$ hours was used leading to 36 degrees of freedom per optimization. On an average, each optimization took 3.6 seconds to compute the optimal control law over a prediction horizon, making the method feasible in real time.

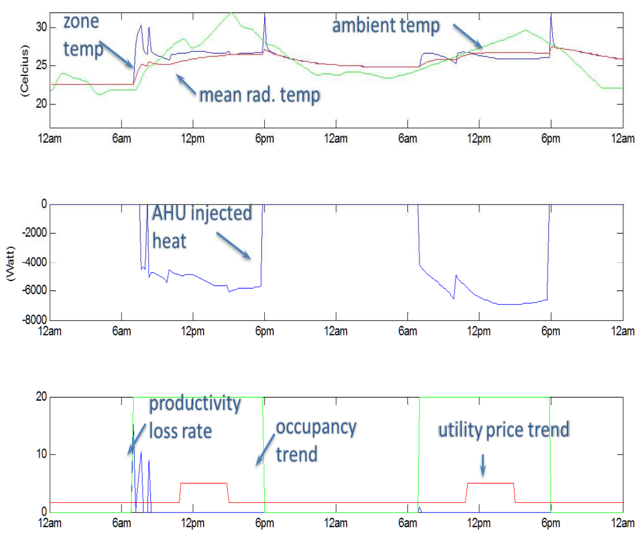

Fig. 5: Trajectories under sequential quadratic programming based solver.

V. Comparative Analysis

Table I summarizes the performance of the MPC controllers along with the conventional controller. We list the salient observations below.

- Both MPC based controllers offer savings in energy costs when compared to conventional control. This is expected as there is provision for taking time-of-day pricing into account in MPC based solvers.
- The AQR based solver offers savings in discomfort costs as well. For the case-study, it is therefore possible to reduce operational costs using MPC based controller. More work is required to quantify the effect of mode switching in the AQR based solver.
- For numerical optimization based SQP solver, convergence to the optimal solution is a major issue. Convergence to a local minima and improper initial guess can result in suboptimal costs. More investigation is needed to identify the reason for the suboptimal performance in discomfort metric.
- In all simulations, it is observed that the discomfort cost is dominated by the spike at the beginning of the occupancy periods (in all simulations). The effect of sampling the discontinuous occupancy gives the impression of high productivity loss on arrival though the notion of productivity at the instant of arrival is unintuitive. More realistic occupancy (and internal gain) profiles are required to evaluate the discomfort cost savings of the solvers.
- The energy costs were calculated based on an optimized lookup table for plant operation. This reduced the available degrees of freedom in actual plant operation. With judicious trade off between fan and pump power, it should be possible to attain more energy savings in practice than those reported.

	Energy Cost (\$/day)	Discomfort Cost (\$/occupant/day)	Total Cost(\$)	Computation time (per decision)
Conventional	0.287	0.720	14.68	–
AQR	0.284 (1.0%) savings	0.517 (28%) savings	10.92 (25.6%) savings	0.20 sec (1 day horizon)
SQP	0.274 (4.5%) savings	1.11 (54%) expensive	22.47 (53.4%) expensive	3.6 sec (6 hr horizon)

TABLE I: Summary of solver performance

VI. Conclusions

We have described an evaluation procedure of model predictive control strategies for supervisory control of building HVAC systems. Using a case-study to benchmark different strategies, we conclude that there is scope for savings using model predictive control. A direct method to integrate discomfort costs into the cost function using occupant productivity has also been proposed. The case-study described can be part of a preliminary toolkit to compare the efficacy of other algorithms. Additional criteria that can be incorporated for comparing solvers include their sensitivity to inaccurate forecasts and models. Future directions include comparing the strategies in multiple and more realistic scenarios.

References

[1] V. Zavala, J. Wang, S. Leyffer, E. Constantinescu, M. Anitescu, and G. Conzelmann, "Proactive energy management for next-generation building systems," *Proc. IBPSA-USA SimBuild 2010*, 2010.

[2] Y. Ma, F. Borrelli, B. Hencey, B. Coffey, S. Bengea, and P. Haves, "Model predictive control for the operation of building cooling systems," in *American Control Conference (ACC), 2010*. IEEE, 2010, pp. 5106–5111.

[3] D. Kim and J. Braun, "Reduced-order building modeling for application to model-based predictive control," in *Proc. IBPSA-USA SimBuild 2012*, 2012.

[4] ——, "Model-based predictive control for buildings with decoupling and reduced-order modeling," in *Proceedings of 2nd International High Performance Building Conference*, 2012.

[5] P. Roelofsen, "The impact of office environments on employee performance: The design of the workplace as a strategy for productivity enhancement," *Journal of facilities Management*, vol. 1, no. 3, pp. 247–264, 2002.

[6] P. Fanger, "Thermal comfort. analysis and applications in environmental engineering." *Thermal comfort. Analysis and applications in environmental engineering.*, 1970.

[7] J. Maciejowski, *Predictive control: with constraints*. Pearson education, 2002.

[8] R. Cagienard, P. Grieder, E. Kerrigan, and M. Morari, "Move blocking strategies in receding horizon control," *Journal of Process Control*, vol. 17, no. 6, pp. 563–570, 2007.