2011 50th IEEE Conference on Decision and Control and
European Control Conference (CDC-ECC)
Orlando, FL, USA, December 12-15, 2011

Analysis of Local Optima in Predictive Control for Energy Efficient Buildings

Anthony Kelman*, Yudong Ma*, Francesco Borrelli*

Abstract— **We study the problem of heating, ventilation, and air conditioning (HVAC) control in a typical commercial building. We propose a model predictive control (MPC) approach which minimizes energy use while satisfying occupant comfort and actuator constraints by using predictive knowledge of weather and occupancy.**

The objective of this paper is to investigate the phenomenon of local optima. In particular, the product between air temperatures and mass flow rates arising from energy balance equations leads to a non-convex MPC problem. Fast computational techniques for solving non-convex optimization can only provide certificates of local optimality.

In the first part of the paper, models and MPC control design for two common HVAC configurations are introduced. In the second part, simulation results exhibiting local optima for both configurations are presented. We perform a detailed analysis on the different types of local optima and their physical interpretation. Our ultimate goal is to use this analysis to derive branch and bound rules which allow a nonlinear programming solver to converge to globally optimal control sequences.

I. Introduction

The building sector consumes about 40% of the energy used in the United States and is responsible for nearly 40% of greenhouse gas emissions [8]. It is therefore economically, socially and environmentally significant to reduce the energy consumption of buildings. Achieving this goal requires the development of highly efficient heating and cooling systems, which are more challenging to control than conventional systems [1], [4], [5].

This work focuses on model predictive control (MPC) of heating, ventilation, and air conditioning (HVAC) over networks of thermal zones. We consider two common configurations of HVAC systems. The first configuration is known as dual-duct, single fan [2]. The second configuration is known as single-duct variable air volume (VAV) with reheat. These configurations are described in detail in Section II.

Predictive controllers are designed for each of the HVAC configurations. The resulting optimization problems are generally non-convex because the system dynamics are nonlinear. Nonlinear programming (NLP) solvers based on sequential quadratic programming [9] or interior point methods [11] cannot guarantee global optimality. In this work, we study the local optima of the MPC optimization problem for the two aforementioned HVAC configurations. We make use of existing NLP solvers in order to handle the system nonlinearities. The set of locally optimal solutions is thoroughly explored by repeatedly running a NLP solver using a large number of initial guesses, distributed throughout the optimization space. Physical explanations are presented for each local optimum. The ultimate goal is to use this analysis to develop tailored branch and bound algorithms with guarantees of convergence to global optimality.

We remark that the evaluation of optimal controllers for building climate regulation has been studied in the past by several authors (see [3], [6], [10] and references therein). Compared to existing literature this paper focuses on the specific issue of global vs local optima of MPC optimization problems for nonlinear systems.

* A. Kelman, Y. Ma, F. Borrelli are with the Department of Mechanical Engineering, University of California, Berkeley, CA 94720-1740. E-mail: {kelman, myd07, fborrelli}@berkeley.edu.

This work has been partially supported by United Technologies Research Center (UTRC) under grant W912-09-C-0056 and by U.S. Air Force Office of Scientific Research (AFOSR) under grant FA9550-09-1-0106.

II. Configuration Descriptions

This section describes the two HVAC system configurations we focus on.

A. Configuration A

The HVAC system configuration known as dual-duct, single fan is shown in Figure 1. Supply air is heated and cooled to desired temperatures in two separate duct systems by a pair of coils (air-to-water heat exchangers). The two separate duct systems route hot and cold air to mixing boxes at each thermal zone (typically one or several rooms). A mixing box contains a pair of linked dampers (position-controlled louvers), designed so that when the hot side damper closes, the cold side damper opens and vice-versa. We refer to these as the *zone dampers*. The zone dampers serve as a control actuator to provide the desired mixed supply air temperature to a zone. The supply air can be set anywhere in the range between the cooling coil outlet temperature and the heating coil outlet temperature.

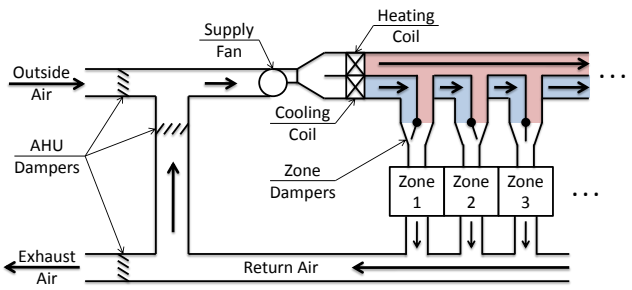

Fig. 1. Dual-duct, single fan HVAC system schematic

Mixed zone air returns to the central air handling unit (AHU) through a return duct. A set of AHU dampers can either exhaust the return air to ambient and use fresh outdoor air as input to the supply fan, or recirculate the return air, or some combination of the two.

We consider systems where the supply fan is equipped with a variable frequency drive (VFD), so the fan speed and therefore the total flow rate can be controlled. Since the mixing box dampers for each zone are linked, the individual zone flow rates are not independently controllable. As a first-order approximation, we treat the flow splits to each zone as constant. This system will be denoted Configuration A in the remaining sections.

In summary, the control inputs in this system are: the total supply fan flow rate, the outdoor air flow rate into the AHU, the cooling coil outlet temperature setpoint, the heating coil outlet temperature setpoint, and the zone mixing box supply temperature setpoints. The states in this system are the zone temperatures.

B. Configuration B

In this section we introduce an alternative HVAC system configuration, known as single-duct variable air volume (VAV) with reheat. We consider an air handling unit serving multiple zones, as before. The AHU in this configuration is capable of using either recirculated zone exit air, fresh outside air, or a mix of the two. As shown in Figure 2, all of the supply air flows through a cooling coil. The cool air is distributed by a fan to the "VAV boxes" at each zone. A VAV box consists of a damper and a heating coil. The damper position in a VAV box controls the flow rate of air supplied to an individual zone. The heating coil is used to warm the supply air if that zone requires heating.

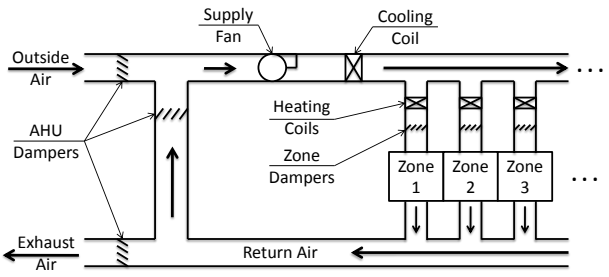

Fig. 2. Single-duct variable air volume with reheat HVAC system schematic

Compared to Configuration A, this single-duct system has local control over heating rather than one central heating coil, and the VAV dampers provide individual control over the supply flow splits delivered to each zone.

In summary, the control inputs in this system are: the flow rates of air supplied to each zone, the outdoor air flow rate into the AHU, the cooling coil outlet temperature setpoint, and the heating coil outlet temperature setpoints at each zone. The states in this system are the zone temperatures.

III. System Modeling

In order to develop control-oriented thermal models of limited order and reduced complexity, we make the following assumptions:

- A1 The average air temperature dynamics of the thermal zones can be reasonably approximated as first-order. We therefore combine the thermal capacitance of the air, walls, furnishings, and other contents of zone i into a single lumped parameter denoted $(mc)_i$.
- A2 Humidity is not explicitly included in our model.
- A3 All dynamics except those of the thermal zones are neglected. Actuators are assumed to instantly meet their control setpoints.
- A4 A prediction of the thermal loads $\dot{Q}_i$ in each zone due to occupants, equipment, and all heat transfer to or from ambient and other zones is known in advance. Predicted outside ambient temperature T_{oa} is also known.

A. Thermal Zone Model

A first order energy balance gives the following continuous time system dynamics for the temperature T_{zi} of zone i

$$(mc)_i \frac{d}{dt} T_{zi} = \dot{Q}_i + \dot{m}_{zi} c_p (T_{si} - T_{zi}), \tag{1}$$

where c_p is the specific heat capacity of air and T_{si} is the temperature of the supply air delivered to zone i. The flow rate $\dot{m}_{zi}$ and supply temperature T_{si} are inputs to this model from the HVAC system.

If we neglect heat transfer by radiation, the thermal loads $\dot{Q}_i$ can be represented as an affine function of the zone temperatures

$$\dot{Q} = \begin{bmatrix} \dot{Q}_1 \\ \vdots \\ \dot{Q}_n \end{bmatrix} = R \begin{bmatrix} T_{z1} \\ \vdots \\ T_{zn} \end{bmatrix} + \dot{Q}_{\text{offset}},$$

where n is the number of thermal zones served by the same AHU, R is a symmetric $n \times n$ matrix of heat transfer coefficients and $\dot{Q}_{\text{offset}}$ is a $n \times 1$ vector of constant terms from the predicted thermal load.

The compact form of (1) for all n zones together is

$$M \frac{d}{dt} T_z = R\, T_z + \dot{Q}_{\text{offset}} + c_p \,\text{diag}(\dot{m}_z)(T_s - T_z), \tag{2}$$

where $M = \text{diag}((mc)_1, \ldots, (mc)_n)$, $T_z = [T_{z1}, \ldots, T_{zn}]^T$, $\dot{m}_z = [\dot{m}_{z1}, \ldots, \dot{m}_{zn}]^T$, and $T_s = [T_{s1}, \ldots, T_{sn}]^T$.

Let $u_z = \{\dot{m}_z, T_s\}$, $A(u_z) = M^{-1}(R - c_p \,\text{diag}(\dot{m}_z))$, $B(u_z) = c_p M^{-1} \,\text{diag}(\dot{m}_z) T_s$, and $w = M^{-1}\dot{Q}_{\text{offset}}$, then (2) has the following state-affine form

$$\frac{d}{dt} T_z = A(u_z) T_z + B(u_z) + w.$$

Assuming $\dot{m}_z$ and T_s are zero-order held at sample rate Δt, we discretize this using the trapezoidal method

$$\frac{T_z^+ - T_z}{\Delta t} = A(u_z) \frac{T_z^+ + T_z}{2} + B(u_z) + \frac{w^+ + w}{2}, \tag{3}$$

where T_z^+ and w^+ denote the corresponding values at the next discrete time step $t+\Delta t$. We choose the trapezoidal discretization as a compromise between simplicity and stability (which is critical since Δt may be large).

This zone thermal model is used for both HVAC configurations, however the mass flow rates $\dot{m}_z$ and temperatures T_s of supply air differ depending on the HVAC configuration.

B. HVAC System Model

1) Configuration A: The HVAC control inputs in this configuration are: the total mass flow rate at the supply fan $\dot{m}_s$, the outdoor air flow rate into the AHU $\dot{m}_{oa}$, the cooling coil setpoint T_c, the heating coil setpoint T_h, and the zone mixing box supply temperature setpoints $T_s = [T_{s1}, \ldots, T_{sn}]^T$.

The portion of the total supply flow delivered to each zone, denoted φ_i for zone i, is constant so that

$$\dot{m}_{zi} = \varphi_i \dot{m}_s. \tag{4}$$

In configuration A, only part of the supply flow passes through the heating and cooling coils. The partial flow rates depend on the zone supply temperatures and flow splits.

$$\dot{m}_h = \sum_{i=1}^{n} \frac{\dot{m}_{zi}(T_{si} - T_c)}{T_h - T_c}, \text{ and } \dot{m}_c = \dot{m}_s - \dot{m}_h, \tag{5}$$

where $\dot{m}_h$ and $\dot{m}_c$ are respectively the mass flow rates through the heating and cooling coils. The coil flows do not directly influence the zone temperature dynamics, but will be important for the coil energy consumption in Section III-C.

2) Configuration B: The HVAC control inputs in this configuration are: the mass flow rates of air supplied to each zone $\dot{m}_z = [\dot{m}_{z1}, \ldots, \dot{m}_{zn}]^T$, the outdoor air flow rate into the AHU $\dot{m}_{oa}$, the cooling coil setpoint T_c, and the heating coil setpoints at each zone $T_h = [T_{h1}, \ldots, T_{hn}]^T$. The supply air temperature T_{si} delivered to zone i is equal to the VAV box heating coil setpoint, $T_{si} = T_{hi}$.

All of the supply flow passes through the cooling coil in configuration B, but only the flow to a single zone passes through each heating coil. The coil flow rates are then

$$\dot{m}_{hi} = \dot{m}_{zi}, \text{ and } \dot{m}_c = \dot{m}_s = \sum_{i=1}^{n} \dot{m}_{zi}, \tag{6}$$

where $\dot{m}_{hi}$ is the flow rate through the heating coil at zone i, and $\dot{m}_c$ is the flow rate through the cooling coil.

3) Common Aspects: In both configurations, we assume the ratio of supply flow rate to return flow rate is the same for all zones. The return air temperature T_r is therefore given by a flow-rate-weighted average

$$T_r = \frac{\sum_{i=1}^{n}(\dot{m}_{zi} T_{zi})}{\sum_{i=1}^{n} \dot{m}_{zi}}. \tag{7}$$

The AHU mixed air temperature T_m is similarly a flow-weighted average of outdoor air temperature T_{oa} and return temperature T_r

$$T_m = \frac{\dot{m}_{oa} T_{oa} + (\dot{m}_s - \dot{m}_{oa}) T_r}{\dot{m}_s}. \tag{8}$$

In our model the return and mixed air temperatures influence the coil energy consumption, but not the zone dynamics.

C. Cost Function: Energy Consumption and Price

Energy use of the cooling and heating coils is calculated as air-side thermal power $\dot{m}\, c_p \Delta T$, based on the models in the previous section. We represent the operating characteristics of the cold and hot water circuits with two parameters: efficiency η_h for the hot side, and coefficient of performance η_c for the cold side. The two HVAC configurations A and B have different flow rates and delta temperatures across the coils, so the expressions for the power used by the coils depend on the configuration.

For both configurations cooling coil power P_c has the form

$$P_c = \frac{c_p}{\eta_c} \dot{m}_c (T_m - T_c), \tag{9}$$

where $\dot{m}_c$ is given by (5) for configuration A and by (6) for configuration B.

For configuration A the heating coil power P_h is

$$P_h = \frac{c_p}{\eta_h} \dot{m}_h (T_h - T_m), \tag{10}$$

where $\dot{m}_h$ is given by (5).

For configuration B the total power of all heating coils is

$$P_h = \sum_{i=1}^{n} \left(\frac{c_p}{\eta_h} \dot{m}_{zi} (T_{si} - T_c) \right). \tag{11}$$

In both configurations the electrical power P_f used by the supply fan is

$$P_f = \frac{\dot{m}_s \, \Delta p}{\rho \, \eta_f}, \tag{12}$$

where Δp is the pressure difference across the fan, ρ is the air density, and η_f is the efficiency of the fan. Assuming incompressible flow gives $\Delta p \propto \dot{m}_s^2$, where the ratio of proportionality depends on the flow resistance of all the downstream zone dampers.

In configuration A we assume the flow resistance of the mixing boxes is constant, so we can take $\Delta p = \rho \, \eta_f \kappa_A \dot{m}_s^2$ where the parameter κ_A captures the fan efficiency and duct pressure losses. Equation (12) then becomes

$$P_f = \kappa_A \dot{m}_s^3. \tag{13}$$

In configuration B however, the flow resistance of the VAV dampers depends on their positions. At higher flow rates with the dampers more open, the overall flow resistance is lower. So the increase of pressure drop with flow rate will be slower than quadratic, and fan power increases slower than cubic. For simplicity we restrict our model to polynomial form, so we take $\Delta p = \rho \, \eta_f \kappa_B \dot{m}_s$. With this form of simplification for configuration B, we have

$$P_f = \kappa_B \dot{m}_s^2. \tag{14}$$

We introduce several parameters to reflect utility pricing. The cost in dollars per unit energy content is denoted r_e for electricity, r_h for heating fuel (typically gas, or steam from

a central plant). The total utility cost from time t to time $t + N\Delta t$, where N is the prediction horizon length, is

$$J = \int_t^{t+N\Delta t} (r_e P_f + r_e P_c + r_h P_h) d\tau. \quad (15)$$

D. Constraints

The system states and control inputs are subject to constraints due to control requirements and actuator limits. Since the control inputs differ for the two HVAC configurations, some of these constraints only apply to one configuration.

1) Configuration A:

- $T_h \geq \underline{T}_m$, heating coil can only increase temperature.
- $T_h \leq \overline{T}_h$, heating coil capacity (hot water temperature).
- $T_{si} \leq T_h \ \forall \ i \in \{1, \ldots, n\}$, hottest supply temperature.

2) Configuration B:

- $\underline{\dot{m}}_{zi} \leq \dot{m}_{zi} \leq \overline{\dot{m}}_{zi} \ \forall \ i \in \{1, \ldots, n\}$, minimum ventilation requirement and maximum VAV box capacity.
- $T_{si} \leq \overline{T}_h \ \forall \ i \in \{1, \ldots, n\}$, heating coil capacity.

3) Common Constraints:

- $\underline{\dot{m}}_s \leq \dot{m}_s \leq \overline{\dot{m}}_s$, minimum overall ventilation requirement and maximum fan capacity.
- $\underline{\dot{m}}_{oa} \leq \dot{m}_{oa} \leq \dot{m}_s$, minimum set by required fresh air for indoor air quality, maximum is entirely fresh air.
- $T_c \leq T_m$, cooling coil can only decrease temperature.
- $T_c \geq \underline{T}_c$, cooling coil capacity (cold water temperature).
- $T_{si} \geq T_c \ \forall \ i \in \{1, \ldots, n\}$, coldest supply temperature.
- $\underline{T}_{zi} \leq T_{zi} \leq \overline{T}_{zi} \ \forall \ i \in \{1, \ldots, n\}$, comfort range.

E. Model Summary

Combining the HVAC system model from Section III-B and the discretized thermal zone model from Section III-A, the consolidated model can be expressed as

$$f(x_{k+1|t}, x_{k|t}, u_{k|t}, w_{k+1|t}, w_{k|t}) = 0. \quad (16a)$$

where $x_{k|t}$ is the value of T_z at time $t + k\Delta t$ predicted at time t, $u_{k|t}$ is the value of all control inputs at time $t + k\Delta t$ predicted at time t, and $w_{k|t}$ is the value of the disturbance inputs $\dot{Q}_{\text{offset}}$ and T_{oa} at time $t + k\Delta t$ predicted at time t.

The constraints from Section III-D can be expressed as

$$g(x_{k+1|t}, x_{k|t}, u_{k|t}, w_{k+1|t}, w_{k|t}) \leq 0. \quad (16b)$$

Define the continuous step cost as

$$J^c_{k|t} = \int_{t+k\Delta t}^{t+(k+1)\Delta t} (r_e P_f + r_e P_c + r_h P_h) d\tau. \quad (16c)$$

The integral is approximated according to the trapezoidal discretization for consistency with the discretization of the state dynamics (3). Let $J_{k|t}$ be the discretization of (16c).

The system model f, control inputs u, cost J, and constraints g are different for the two HVAC configurations.

For configuration A, the model f applies the zone supply temperatures and flows, with flow splits from (4), to the thermal zone model (3). The control inputs are $u = \{\dot{m}_s, \dot{m}_{oa}, T_c, T_h, T_s\}$. The cost function J combines (5), (7)-(10), and (13). The constraint function g combines Section III-D.1, Section III-D.3, (7), and (8).

For configuration B, the model f applies the zone supply temperatures (with $T_s = T_h$) and flows to the thermal zone model (3). The control inputs are $u = \{\dot{m}_z, \dot{m}_{oa}, T_c, T_h\}$. The cost function J combines (6)-(9), (11), and (14). The constraint function g combines Section III-D.2, Section III-D.3, (7), and (8).

IV. Control Design

Model predictive control solves at each time step t the following optimization problem

$$\begin{aligned} &\min_{\mathbf{U},\mathbf{X}} \sum_{k=0}^{N-1} J_{k|t} \quad (17) \\ &\text{subj. to, } \ \forall \ k \in \{0, \ldots, N-1\}, \\ &f(x_{k+1|t}, x_{k|t}, u_{k|t}, w_{k+1|t}, w_{k|t}) = 0 \\ &g(x_{k+1|t}, x_{k|t}, u_{k|t}, w_{k+1|t}, w_{k|t}) \leq 0 \\ &x_{0|t} = T_z(t) \end{aligned}$$

where $\mathbf{U} = \{u_{0|t}, \ldots, u_{N-1|t}\}$ is the set of predicted control inputs at time t, $\mathbf{X} = \{x_{1|t}, \ldots, x_{N|t}\}$ is the set of predicted system states at time t, starting from initial state $x_{0|t} = T_z(t)$ and applying the input sequence $\mathbf{U}$ to the system model (16a).

Let the optimal solution of problem (17) at time t be denoted by $\mathbf{U}^\star = \{u^\star_{0|t}, \ldots, u^\star_{N-1|t}\}$. Then, the first step of $\mathbf{U}^\star$ is input to the system, $u(t) = u^\star_{0|t}$. The optimization (17) is repeated at time $t + \Delta t$, with the updated new state $x_{0|t+\Delta t} = T_z(t+\Delta t)$ yielding a *moving or receding horizon control* strategy.

The optimization problem (17) has nonlinear cost and nonlinear constraints. In order to solve this optimization problem we use the interior-point NLP solver Ipopt [11] via the YALMIP [7] toolbox. Ipopt and most other NLP codes are generally not global solvers for non-convex problems, so these algorithms can converge to a local optimizer.

V. Local Optima Analysis

In this section we show simulation data exhibiting local optima for HVAC configurations A and B. The local optima are found by repeated execution of a NLP algorithm. Each execution of the NLP algorithm is started from a randomly selected initial guess, uniformly distributed within the allowable range of states and control inputs. Clearly this method is not guaranteed to find every single local optimum point. However it is expected that over many samples we will find the local optima with the largest regions of attraction.

We take simple instances of the models (16) with three zones and a prediction horizon of 2 steps. For full details regarding paramater values used in this section and complete tabular simulation results, see the technical report available at `www.mpc.berkeley.edu/people/tony-kelman`.

A. Configuration A

Our simulations revealed the presence of local optima for configuration A in a scenario with small zones (low thermal capacitance) subject to large loads. We found 6 distinct families of local optima in this scenario. Local optima within

the same family have the same cost, and those belonging to different families have different associated optimal costs. The six families can be further explained as follows.

At the first time step there are two different control modes: a heating mode and a cooling mode. At the second time step, there are three different control modes: a heating mode, a cooling mode, and a third "intermediate" mode. Every combination of modes for the first and second time step was feasible, so over the horizon of 2 steps we have six families of local optima. Points belonging to the same family have different heating or cooling coil setpoints. The cost value, states, and all other control inputs are *equal* within a family. The families are illustrated in Figure 3 and the corresponding cost values are given in Table I.

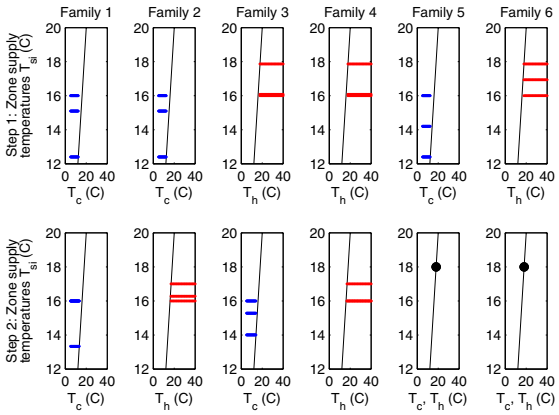

Fig. 3. Families of local optima for HVAC configuration A.

TABLE I
COST VALUES OF LOCAL OPTIMA, CONFIGURATION A

Family #	1	2	3	4	5	6
$J^\star$ ($)	0.095	0.113	0.133	0.143	0.177	0.21

In this scenario the zones have positive thermal loads $\dot{Q}_i$ so the supply temperatures must be lower than the zone temperatures in order to counteract the loads and remain within the comfort range. The ambient temperature here is $T_{oa} = 16$ °C, which is cooler than the zone temperatures. Free cooling is available by using outside air instead of recirculated return air at the AHU, reducing coil energy.

In the *cooling mode* the cooling coil setpoint $T_c < 16$ °C and the heating coil setpoint $T_h = 16$ °C (no heating of the outside air). In the *heating mode* the heating coil setpoint $T_h > 16$ °C and the cooling coil setpoint $T_c = 16$ °C (no cooling of the outside air). The *intermediate mode* is a special case where a single supply temperature can meet all of the different thermal loads. That specific supply temperature for this case is 18 °C which is warmer than ambient temperature, so this mode can be realized by mixing ambient air with recirculated air. When AHU mixed temperature $T_m = 18$ °C, the cooling and heating coil setpoints can both be set to 18 °C requiring no coil energy at all. This mode requires zones with different loads to be at different temperatures, so it is not possible at the first time step when all of the zones have the same initial condition temperature.

Note that in cooling mode, the cooling coil setpoint T_c can have any value between 5 and 16 °C as long as it is cooler than the lowest required zone supply temperature. Likewise in heating mode, the heating coil setpoint T_h can have any value between 16 and 40 °C as long as it is warmer than the highest required zone supply temperature. When T_c is cooler or T_h warmer than necessary, the zone dampers compensate to maintain the same mixed supply temperatures T_s.

Cooling mode is feasible for low supply flow rates, where all required zone supply temperatures are below 16 °C (otherwise the heating coil would also need to be active). Heating mode is feasible at higher flow rates, where all required supply temperatures are above 16 °C. The intermediate mode is only feasible for very specific zone temperatures, and had the highest flow rate of the three modes here.

Table I confirms that the modes requiring lower flow rates have lower overall cost (accounting for coil energy as well). The lowest-cost family of local optima is in cooling mode at both time steps. Due to the initial conditions, heating mode at the second time step requires less flow than heating mode at the first time step to remain inside the comfort bounds. So the second-lowest-cost family of local optima is in cooling mode at the first step, then heating mode at the second step. The third-lowest-cost family is the reverse: heating mode then cooling mode. The fourth-lowest cost family is in heating mode at both steps. The fifth-lowest cost family is in cooling mode then the intermediate mode. The highest-cost family of local optima is in heating mode then the intermediate mode.

B. Configuration B

Our simulations revealed the presence of local optima for configuration B in a scenario with large zones (high thermal capacitance) subject to small loads and a time-varying comfort bound. We found 6 distinct local optima for configuration B in this scenario. Each of the local optima here was a single point, rather than a connected family of solutions as in configuration A.

For all six local optima, only outdoor air is used so $\dot{m}_s = \dot{m}_{oa}$, and the heating and cooling coils are inactive with $T_c = T_h = 16$ °C at all times. In this scenario outside air is sufficient to counteract the zone loads and cool the zones to remain within the time-varying comfort bounds.

The six optima are qualitatively very similar to each other: in the first time step, one zone is cooled to slightly below 24 °C, another zone is cooled to approximately 25 °C, and the third zone remains at the initial upper bound of 26 °C (see the upper row of Figure 4). In the second time step, all 3 zones are controlled to reach the reduced upper bound at 24 °C. Mass flow rates are large for time steps when a zone requires a 2 °C temperature change, intermediate for steps when a zone requires a 1 °C temperature change, and small for steps when the zone temperature change is small (see the lower row of Figure 4). The difference between the local optima here is a matter of sequencing: the local optima correspond to the 6 different permutations of 3 zones.

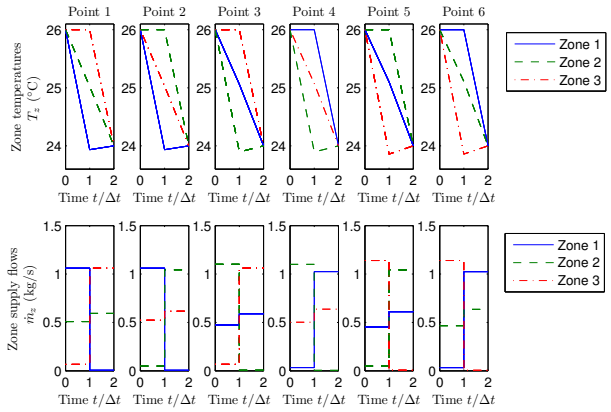

Fig. 4. Local optima for HVAC configuration B.

TABLE II
COST VALUES OF LOCAL OPTIMA, CONFIGURATION B

Point #	1	2	3	4	5	6
$J^\star$ ($)	0.11738	0.11751	0.11753	0.11779	0.11783	0.11796

Table II gives the corresponding cost values for the 6 locally optimal points. The cost values show very little variation between the local optima, only 0.5% difference between the lowest cost and highest cost points. The differences in cost are due to the thermal load values for each zone.

For a constant supply temperature (16 °C), less flow is required to counteract a positive thermal load when zone temperature is higher due to (1). So the control strategy that cools the zones with the lowest thermal loads first is the most efficient here, keeping the high-load zones at higher temperatures for longer. This strategy is point 1, so the results confirm this statement. Point 2 partially cools the highest-load zone before the intermediate-load zone, and point 3 fully cools the intermediate-load zone before the lowest-load zone. Point 4 fully cools the lowest-load zone last instead of first but is otherwise in order, and point 5 fully cools the highest-load zone first instead of last but is otherwise in order. Point 6 is the reversed order from point 1.

The small difference in cost between points 2 and 3, and between points 4 and 5, is due to which zone is fully cooled at the first time step. That zone (zone 1 for points 1 and 2, zone 2 for points 3 and 4, zone 3 for points 5 and 6) then uses exactly the lower bound minimum flow rate at the second time step. At the minimum flow rate, the change in zone temperature during the second time step is larger for higher thermal loads. Therefore the higher the thermal load, the more a zone must be overcooled at the first time step in order to remain below 24 °C at the second time step. This additional energy is another factor causing points 5 and 6 to have higher cost than points 3 and 4, and points 3 and 4 to have higher cost than points 1 and 2.

VI. CONCLUSIONS

We have studied the problem of using model predictive control (MPC) for heating, ventilation, and air conditioning (HVAC) systems in a typical commercial building for two common HVAC configurations. We investigated the phenomenon of local optima by solving the nonlinear MPC problem using the local nonlinear programming solver Ipopt. We have shown by simulation that the resulting optimization problem exhibits local optima for both configurations. We performed a detailed analysis of the different types of local optima and their physical interpretation.

For HVAC configuration A, the local optima had notably different energy costs, so finding a globally optimal control strategy has obvious significance. In configuration B, the local optima were very similar in cost for our particular scenario. We are working to identify which factors can make the differences more pronounced. Even the presence of local optima with practically equivalent costs but significantly different control strategies could have important consequences. For example, switching between local optima at successive time steps could cause oscillatory control inputs, which is typically undesirable and detrimental to actuator reliability.

Our future research will focus on applying the type of analysis presented in this work to a generic scenario (in terms of number of zones, length of horizon, system parameters, and operating conditions). We expect the number of local optima to grow for models with more zones over longer horizons, which will exacerbate the issue. Our ultimate goal is to use the proposed analysis in order to derive branch and bound rules which allow a NLP solver to converge to globally optimal control sequences.

REFERENCES

[1] D.B. Crawley, L.K. Lawrie, F.C. Winkelmann, W.F. Buhl, Y.J. Huang, C.O. Pedersen, R.K. Strand, R.J. Liesen, D.E. Fisher, M.J. Witte, and J. Glazer. Energyplus: creating a new-generation building energy simulation program. *Energy and Buildings*, 33(4):319 – 331, 2001.

[2] W.T. Grondzik. *Air-conditioning system design manual*. ASHRAE professional series. Butterworth-Heinemann, 2007.

[3] G.P. Henze, C. Felsmann, and G. Knabe. Evaluation of optimal control for active and passive building thermal storage. *International Journal of Thermal Sciences*, 43(2):173 – 183, 2004.

[4] M. Koschenz and V. Dorer. Interaction of an air system with concrete core conditioning. *Energy and Buildings*, 30(2):139–145, 1999.

[5] M. Koschenz and B. Lehmann. Development of a thermally activated ceiling panel with PCM for application in lightweight and retrofitted buildings. *Energy and Buildings*, 36(6):567–578, 2004.

[6] S. Liu and G.P. Henze. Experimental analysis of simulated reinforcement learning control for active and passive building thermal storage inventory: Part 1. theoretical foundation. *Energy and Buildings*, 38(2):142–147, 2006.

[7] J. Löfberg. Yalmip : A toolbox for modeling and optimization in MATLAB. In *Proceedings of the CACSD Conference*, Taipei, Taiwan, 2004.

[8] J.M. McQuade. A systems approach to high performance buildings. Technical report, United Technologies Corporation, Apr 2009. `http://gop.science.house.gov/Media/hearings/energy09/april28/mcquade.pdf`.

[9] J. Nocedal and S. J. Wright. *Numerical Optimization*, chapter 18. Springer-Verlag, 2006.

[10] F. Oldewurtel, A. Parisio, C.N. Jones, M. Morari, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, and K. Wirth. Energy efficient building climate control using stochastic model predictive control and weather predictions. In *2010 American Control Conference*, pages 5100–5105, Jun 2010.

[11] A. Wächter and L. T. Biegler. On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106(1):25–57, 2006.