Energy and Buildings 82 (2014) 675–684

Contents lists available at ScienceDirect

Energy and Buildings

journal homepage: www.elsevier.com/locate/enbuild

Model predictive HVAC control with online occupancy model

Justin R. Dobbs *, Brandon M. Hencey

Department of Mechanical and Aerospace Engineering, Cornell University, Upson Hall, Ithaca, NY 14853, USA

ARTICLE INFO

Article history:
Received 3 March 2014
Received in revised form 7 July 2014
Accepted 22 July 2014
Available online 1 August 2014

Keywords:
Model predictive control
MPC
Occupancy prediction
On-line training
Markov chains
HVAC

ABSTRACT

This paper presents an occupancy-predicting control algorithm for heating, ventilation, and air conditioning (HVAC) systems in buildings. It incorporates the building's thermal properties, local weather predictions, and a self-tuning stochastic occupancy model to reduce energy consumption while maintaining occupant comfort. Contrasting with existing approaches, the occupancy model requires no manual training and adapts to changes in occupancy patterns during operation. A prediction-weighted cost function provides conditioning of thermal zones before occupancy begins and reduces system output before occupancy ends. Simulation results with real-world occupancy data demonstrate the algorithm's effectiveness.

1. Introduction

The long-term increase in energy prices has driven greater interest in demand-based HVAC control. Fixed temperature setpoint schedules and occupancy-triggered operation are commonly used to trim energy consumption, but these approaches have significant drawbacks. First, fixed schedules become outdated; when occupancy patterns change, early or late occupants are left uncomfortable, or the space is conditioned prematurely or for too long. Second, thermal lag limits response speed and thus precludes aggressive temperature set-back. Addressing both schedule inaccuracy and thermal lag requires a stochastic occupancy model and a control scheme that can use it effectively.

Considerable research effort has been directed toward occupancy detection and modeling. Work on detection has focused on boosting accuracy through sensor fusion using probabilistic, neural, or utility networks [1–4]. Agent-based models have been used to predict movement within buildings [5,6], as have Markov chains [7–9]. Erickson and Dong, for example, considered rooms to be Markov states and movements among them to be transitions in order to predict persons' behavior, while Dong and Lam [10] used a semi-Markov model to merge multiple sensor streams into an occupant count estimate. The simpler Page model considered Boolean occupancy (occupied or vacant) under a time-heterogeneous Markov chain to generate realistic simulation input data, rather than for on-line forecasting [8].

With the exception of the Page model, the above efforts have found use in heuristic [11–13] or model predictive control (MPC) schemes [10,13,14], but they face barriers to widespread usage. Most notably, where authors have used MPC, they have also used manually-generated thermal models [10,13,14] even though model creation is tedious and time-consuming and therefore expensive. Eager to demonstrate excellent performance, researchers have favored systems with complex topologies and numerous adjustments that yield "one-off" engineering efforts without a clear path to large-scale adoption. The system outlined in [10], for instance, uses CO_2, sound, and light sensors that require carefully set detection thresholds for each room, plus an on-board weather forecasting algorithm in lieu of forecasts already available. We aim, instead, to make occupancy-predicting control accessible to a broader audience by presenting a simple but effective algorithm with a straightforward implementation. For example, we use an automated BIM translation facility outlined in a previous paper [15], and the core algorithm is industry-standard MPC with occupancy weighting in the cost function. Each of the very few adjustments serves a clearly defined purpose, and we have outlined each component's operation with the practitioner in mind.

Second, recent research has paid little attention to the commissioning and maintenance of occupancy prediction algorithms; model training, if mentioned at all, has been assumed to be done all at one time by someone skilled in the art [8,10,11]. Although

* Corresponding author. Tel.: +1 607 269 5352.
E-mail addresses: jrd288@cornell.edu (J.R. Dobbs), bmh78@cornell.edu (B.M. Hencey).

http://dx.doi.org/10.1016/j.enbuild.2014.07.051
0378-7788/

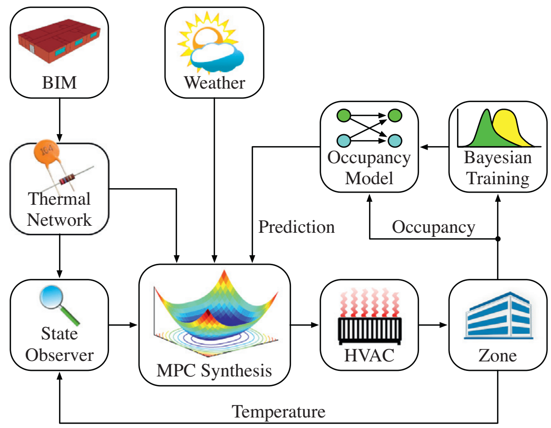

Fig. 1. Proposed system architecture. For this study, the building model has been translated automatically from CAD data into a linear, time-invariant network that encompasses the dominant thermal processes. (Model translation may also be performed manually.)

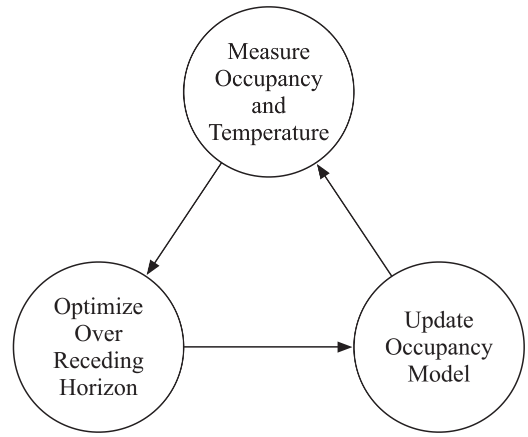

Fig. 2. Process flow during operation.

most training algorithms could be extended to work on-line, ongoing maintenance remains a source of long-term cost neglected by the literature. An occupancy model invariably becomes out-of-date unless it is periodically retrained or can incrementally refine itself with new observations. Our work uses on-line Bayesian inference for stable performance without ongoing manual effort.

The paper progresses as follows. First, we outline the problem formulation. Second, we describe the stochastic occupancy model and its on-line training algorithm. Third, we discuss its integration with model predictive control. Finally, we present simulation results using real-world occupancy data and compare our method's performance to a correctly set scheduled controller and to an occupancy-triggered controller. Throughout the discussion, the control scenario is kept deliberately simple to emphasize the contribution of occupancy learning and its use with MPC.[1]

2. Problem statement

We wish to minimize the total energy usage of a building heating (or cooling) system while maintaining occupant comfort. Versus conventional occupancy-triggered or scheduled control, we aim to

- boost comfort by conditioning the space before occupants arrive,
- limit energy consumption by not running the system too early, and
- exploit stored thermal energy by reducing power before occupants leave.

Our approach is based on MPC but uses a cost function weighted by occupancy predictions from a self-training stochastic model (Fig. 1). At each step, the system measures how much of the previous hour the space was occupied, and the expected occupancy is used to find the best sequence of N future heat inputs to the thermal zone that minimizes the expected cost. The optimization is

$$\begin{aligned} \min_{u_k \cdots u_{k+N-1}} \quad & \sum_{j=0}^{N-1} \mathbb{E}[g(x_{k+j}, u_{k+j}, \tau, \Gamma_{k+j})] \\ \text{subject to} \quad & x_{i+1} = Ax_i + B_u u_i + B_w \mathbb{E}[w_i] \quad \forall i \in \mathbb{Z}^+ \\ & 0 \le u \le u_{\max} \end{aligned} \tag{1}$$

where

- $k \in \mathbb{Z}^+$ is the current time step, and $j \in [0, N-1]$ is the optimization index over the horizon;
- $A \in \mathbb{R}^{n\times n}$ describes the building's thermal dynamics;
- $x \in \mathbb{R}^{n\times 1}$ contains the building's thermal state;
- $u_{k...k+N-1}$ contains the controller output, constrained within the system's capacity $u_{\max}$;
- B_u is a vector that connects the heat input u to the zone air volume;
- w_k is the current weather observation, and $w_{k+1...k+N-1}$ contain an up-to-date weather prediction;
- B_w is a vector that connects the weather conditions to the building envelope;
- τ is the temperature setpoint, which is constant for this study (but can be varied in practice);
- Γ_k is the latest occupancy measurement, and $\Gamma_{k+1...k+N-1}$ are the predicted occupancies; and
- $g(x, u, \tau, \Gamma)$ is a cost function that penalizes total energy consumption and penalizes discomfort based on the occupancy Γ.

The expectation operator $\mathbb{E}[g]$ in Eq. (1) reflects that future values of g require predictions of occupancy and of the weather. The optimization yields a sequence of N power commands to the HVAC system, where positive values are heat and negative are cooling; the first command u_k is applied, and the rest are discarded. The previous and current occupancy observations are then used to train the occupancy model, and the entire process repeats the next time step (Fig. 2).

Two assumptions are made in this presentation. First, we treat the weather forecast as accurate so that we can later omit the expectation operator from w. Second, we use a very simple cost function with constant efficiency and a single linear actuator. These assumptions improve clarity but are not required in practice. Where available, weather uncertainty data can be rolled into the cost function in order to improve robustness [14]. Multiple actuators (e.g. radiant and forced air with vastly different response times) or nonlinear actuation (e.g. variable air volume damper position) can be pulled into the dynamical model and the cost function without undermining the basic approach [13,14,16]. Finally, the energy penalty gain can be varied over time to reflect, for example, changing system efficiency or electricity cost.

3. Building thermal model

Thermal model accuracy influences controller performance, so we need a thermal model that closely approximates the dominant dynamics. Here, we outline how the state-space building model is generated, and we validate it against EnergyPlus simulation results.

[1] See [13] for a comparison of MPC and heuristic control for a more complex HVAC system.